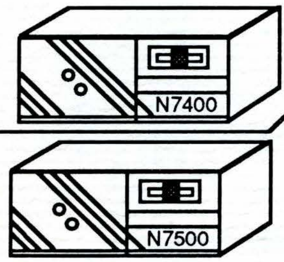

Nucleus 7400/7500

Network Switches and their companion

the Nucleus 7900 *Network Manager*

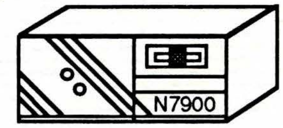

Technical Description Manual

September 1991

AMNET

SECTION 2 ARCHITECTURE & THEORY OF OPERATION (Cont.)

SECTION 3 NUCLEUS 7900 NETWORK MANAGER

SECTION 4 NETWORK SERVICES

SECTION 5 *NETWORK OPERATIONS*

SECTION 6 *HOT STANDBY SWITCH*

SECTION 7 *PRODUCT SPECIFICATIONS*

APPENDIX A *CALL VALIDATION PROCESS*

APPENDIX B *HIGH-SPEED OPTIONS*

APPENDIX C *INTEROPERABILITY*

APPENDIX D *IPL OPERATION*

GLOSSARY

List of Figures

List of Tables

ABSTRACT

Abstract

The abstract covers the structure and the conventions used in the document. This document is designed to assist you in understanding the unique features and capabilities of Release 3.0 on AMNET's Network Switches and Network Manager. It is intended for use by a network administrator, network design engineers, technician, or user's who requires more than just a cursory view of the capabilities of the Nucleus 7400 (N7400) and Nucleus 7500 (N7500) Network Switches, or Nucleus 7900 (N7900) Network Manager.

Release 3.0 Overview

The new features and products associated with Release 3.0 are as follows:

- Nucleus 7500 Network Switch
- Integrated Call Processing
- Nucleus 7900 Network Manager
- T-1/E-1 Line Processor Cards
- Fast Trunk Protocol
- X.25 Facility Enhancements
- Public Data Network Trunk and Gateway
- Backup Trunk
- On-line Port Parameter Changes
- Async PAD Enhancements
- Call Accounting Enhancements

Document Conventions

This type face represents general text. A majority of the documentation is in this text style. The actual tasks or programs that are executable on the node or on the N7900 Network Manager are referenced in lower case *italics*. Disk file names appear in UPPER CASE either on a floppy disk or hard disk. Tasks and system utilities that are an integral part of the Line Processor (LP) cards are listed in *UPPER CASE ITALICS*.

Document Organization

The following outline highlights the general contents of each section in this manual.

SECTION 1	**EXECUTIVE OVERVIEW** Gives an overview of the major components of the Nucleus 7000 Series family. It also provides information on the integrated PAD support, characteristics, and applications supported by the Network Switches.
SECTION 2	**ARCHITECTURE & THEORY OF OPERATION** Provides detailed characteristics and operational techniques implemented in the N7400 and N7500 Network Switches.
SECTION 3	**NUCLEUS 7900 NETWORK MANAGER** Explains the functions unique to the N7900 including the graphics user interface, and the use of dedicated or distributed N7900s.
SECTION 4	**NETWORK SERVICES** Offers the user additional information on virtual circuits, optional user facilities, and inherent network features of the N7400 and N7500.
SECTION 5	**NETWORK OPERATIONS** Provides the operator and the experienced technician with information on network test services, alarms, events, statistics, performance, and reports.
SECTION 6	**HOT STANDBY SWITCH** Presents the product specific features for AMNET's Hot Standby Switch for networking redundancy capabilities.
SECTION 7	**PRODUCT SPECIFICATIONS** Encapsulate the product specific specifications for a quick overview including the Line Processor cards and their options.
APPENDIX A	**CALL VALIDATION PROCESS** Illustrates the call validation process.
APPENDIX B	**HIGH-SPEED OPTIONS** Presents the product specific features for support of the optional high-speed LP cards as well as the IPL operation.
APPENDIX C	**INTEROPERABILITY** Highlights the interoperability of AMNET's products including our N7200 Multiprotocol Switching PAD.
APPENDIX D	**IPL OPERATION** Covers in detail the IPL operation of the Standard Line Processor card.
GLOSSARY	

Abbreviations

The following abbreviations will be encountered throughout this manual. Also, at the end of the manual is a glossary of terms which provides additional explanation of some of the abbreviations listed here and others that you may encounter throughout this manual.

ACK	Acknowledgement
Async	Asynchronous
BPS	Bits Per Second
CCITT	Consultative Committee for International Telephone and Telegraph
CD	Carrier Detect
CEPT	Conference of European Postal and Telecommunications Administrations
CIP	Computer Interface Process
CMM	Common Memory Module
CSU	Channel Service Unit
CTS	Clear To Send
CUG	Closed User Group
D-4	T-1 Framing format compliance with AT&T definitions
DCE	Data Communications Equipment
DISC	Disconnect
DMA	Direct Memory Access
DOD	Department of Defense
DSU	Data Service Unit
DSX-1	Digital Signal Cross-Connect Level 1
DTE	Data Terminal Equipment
DTR	Data Terminal Ready
E-1	AMNET's E-1 Line Processor card operating at 2.048 Mbps
EIA	Electronic Industries Association
ELP	Event Logging Printer
ESF	Extended Superframe Format
FCS	Frame Check Sequence (error condition)
FEPOS	Front End Processor Operating System (developed by AMNET)
HDLC	High Level Data Link Control
HPAD	Host-to-Terminal PAD
IC	Integrated Circuit
ICP	Integrated Call Processing optional software package on the N7500
IP	Internet Protocol
IPL	Initial Program Load
ISDN	Integrated Services Digital Network
ISO	International Standards Organization
LAPB	Link Access Procedures Balanced
LCN	Logical Channel Number
LP	Line Processor card(s)
N7000	See "Nucleus 7000"
N7200	See "Nucleus 7200"
N7400	See "Nucleus 7400"
N7500	See "Nucleus 7500
N7900	See "Nucleus 7900"
NBM	Node Buffer Memory
NMP	See "Nucleus 7900"
Node	Either a Nucleus 7400 or Nucleus 7500
NPAD	Integral Host-to-Host PAD
NSN	Network Switching Nodes includes both the N7400 and N7500

Nucleus 7000	Family of AMNET's Nucleus 7000 Series of products which are the N7200, N7400, N7500, and Nucleus 7900 Network Manger.
Nucleus 7200	AMNET's Multiprotocol Switching PAD
Nucleus 7400	AMNET's Network Switch (DOS based)
Nucleus 7500	AMNET's Network Switch (UNIX based)
Nucleus 7900	The Network Manager used to control, configure, and monitor AMNET's entire family of networking products.
NUI	Network User Identification
OSI	Open Systems Interconnection
PAD	Packet Assembler/Disassembler
PDN	Public Data Network
PIA	Parallel Interface Adapter
Port	Is the entry and/or exit point for node traffic.
PPMOS	Packet Processor Operating System
PPS	Packets Per Second
PVC	Permanent Virtual Circuit
QLLC	Qualified Logical Link Control
REJ	Reject
RNR	Receiver Not Ready
RPOA	Registered Private Operating Agent
RR	Receiver Ready
SABM	Set Asynchronous Balanced Mode
SCC	Serial Communications Controllers (8530 IC or equivalent)
SDLC	Synchronous Data Link Control
SNA	Systems Network Architecture (IBM's de facto standard)
SPF	Shortest Path First algorithm
SVC	Switched Virtual Circuit
T-1	AMNET's T-1 Line Processor card operating at 1.544 Mbps
TCP	Transmission Control Protocol
UA	Unnumbered Acknowledgement
UPS	Uninterruptible Power Source
VC	Virtual Circuit
X.3 protocol	CCITT recommendation that defines the service provided by a PAD to connect a character mode terminal to an X.25 network.
X.25 port	The physical connection in the network that provides the data paths between nodes.
X.25 protocol	CCITT recommendation that specifies the interface between DTE & DCE operating in the packet mode on public data networks.
X.28 protocol	CCITT recommendation that defines the user interface to the X.3 services.
X.29 protocol	CCITT recommendation that defines the use of X.25 packets to carry data between a terminal and a host.
X.121	CCITT recommendation for international internetwork addressing.
XPAD	Transparent PAD

SECTION 1

EXECUTIVE OVERVIEW

Introduction

The Nucleus 7000 Series of networking products are defined in this technical description manual. The members of this robust family of networking solutions includes our Nucleus 7400 and Nucleus 7500 Network Switches. We refer to them as our Network Switching Nodes (NSN) and hence the name for the document - the *Nucleus NSN Technical Description Manual*. A companion product to the NSNs is our Nucleus 7900 Network Manager. The Network Manager is part of AMNET's *Integrated Networking Strategy*.

The NSNs are flexible, dynamic, and expandable to serve data networks, hybrid packet/circuit networks, and large diverse networks. They are optimized for high performance switching applications with an integral Packet Assembly/Disassembly (PAD) and adherence to domestic and international standards. This provides for virtually any computer or user device to communicate on the network. The Nucleus 7000 network solution features:

- Dynamic adaptive routing, to direct traffic over the best available path for the highest level of networking reliability
- Network link speeds up to 2 Mbps
- Sustained throughput of over 1,200 packets per second (pps)
- Packet sizes up to 4,096 bytes, allowing you to connect into any network
- Multiple trunk protocols supported, for sensible use of internodal resources
- Automatic load balancing and congestion control, for efficient use of networking resources

AMNET's NSNs perform two important functions: one, they provide a means of fail-safe data transportation; and two, they allow the integration of diverse applications and/or protocols through a common network facility. Compatibility with the user devices is provided in the NSNs through the integrated PAD facilities. In order to support these functions AMNET developed an integrated software architecture corresponding to the three major hardware subsystems:

- Line Processor software
- Packet Processor software
- Network Management software

The Line Processor software is responsible for the physical Level I and link Level II interfaces. It contains the LAPB (trunk and X.25), the 3X integrated PAD, and the SNA/SDLC PAD applications. It includes an operating system and application programs that are downloaded when the NSN is booted, powered on, re-IPLed, or brought on-line. Each LP card has a dedicated 16-bit processor with over half-a-megabyte of internal on-board memory to support these programs and for packet data storage. The LP cards directly interface into the workstation via the ISA Bus architecture.

The Packet Processor (PPM) software operates in the workstation and is designed to perform the routing, addressing, flow and congestion control, priority determination, performance monitoring, and fault detection and isolation. All of the software modules, tasks, and programs executing in the workstation are controlled by the PPM Operating System (PPMOS). This operating environment provides task dispatching and memory management for these tasks. The software that runs in the workstation is logically divided into a Node Buffer Module (NBM), a Common Memory Module (CMM), and a Packet Processor Module (PPM) all controlled by the PPMOS.

The NSNs are delineated by the fact that the Nucleus 7400 operating environment requires MS-DOS and the Nucleus 7500 uses UNIX™ System V/386. Also, the N7400 supports 264 simultaneous connections while the N7500 can handle 2,000 connections.

Figure 1-1 depicts a desktop unit of our NSN. It is expandable to support up to five LPs (the rack mount version can support up to eight LPs). Each LP card provides four RS-232-C (V.24/V.28) interfaces with an aggregate synchronous data rate of up to 256 Kbps. Asynchronous data rates up to 38.4 Kbps are also supported. All ports are software configurable and support user access as well as trunk protocols. A 4-port asynchronous daughter card is available with our Standard LPs providing you with a total of eight ports per LP for a nodal complement of 64-ports. Figure 1-1 illustrates a basic switching node.

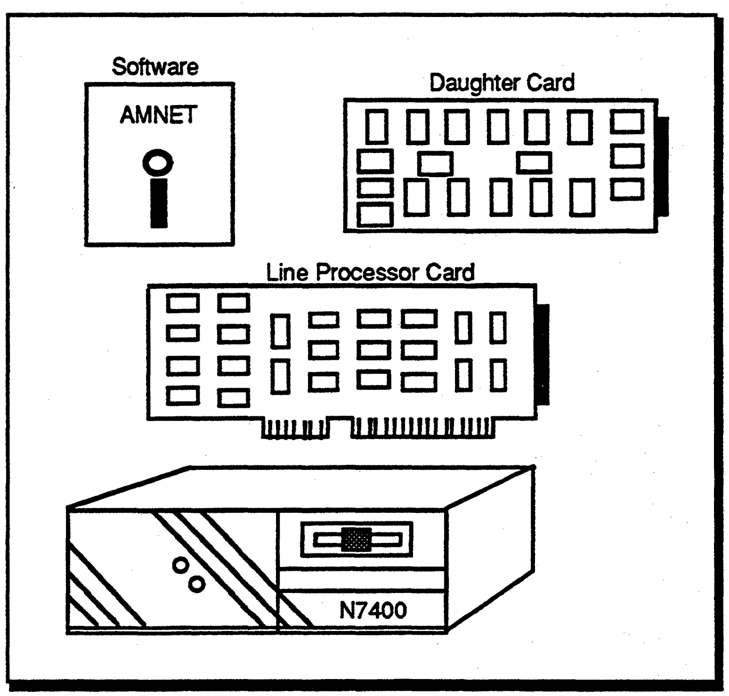

Figure 1-1 Node Architecture

The third major software package consists of a multitasking operating system (UNIX), with graphical user interface, and utility programs designed to aid the network operator and administrator in configuring, monitoring, and controlling the entire network. This is incorporated in our N7900 Network Manager. The various programs and tasks allow the operator to update the database, examine statistical information, and change network parameters while the nodes are operational. Further, it can start up and shut down selected nodes and selectively alter functions on-line at the port level. Its color graphics interface is used in all modes of operation - configuration, monitoring, administration, and database management.

Product Overview

The N7400 and N7500 are members of AMNET's family of high performance Network Switching Nodes. The NSNs provide a modular design for incremental network expansion and are based on industry standards both from a hardware and software perspective. First, the workstation incorporates a powerful 80386 CPU and ISA Bus architecture; second, AMNET's software is written in the high-level 'C' language providing portability; and third, the product adheres to domestic and international standards like CCITT Recommendations X.25 (LAPB), Asynchronous (X.3, X.28, X.29), X.121, and SNA/SDLC.

Both NSNs can be expanded from 4- to 32-ports, using our standard 4-port LP card. Additional PAD expansion is provided with our 4-port Async daughter card. When used in conjunction with our Standard LP cards up to 64-port are supported on a nodal basis. Each port is software selectable on a port-by-port basis. Our high performance LP cards support synchronous and trunk protocols up to 256 Kbps. Higher internodal bandwidth is supported via our T-1 or E-1 LP cards, providing access to high-speed digital services (1.544 Mbps or 2.048 Mbps, respectively). Our NSNs provide an impressive 1,200 pps throughput and configurations are available in both desktop or rack mount units.

A unique dynamic adaptive routing system, directs each packet to its destination via the most efficient path with transparent alternate routing which maximizes throughput and availability. This routing algorithm provides reliable network operations by automatically routing traffic around congested nodes, failed nodes or trunks as soon as a failure is detected in the network topology. In fact, calls in process; as they traverse the network; consult each intermediate node to determine the best trunk over which to forward the packet for optimal performance and reliability.

Default packet window and packet size, flow control and throughput class negotiations are specified on a per port basis. Also, barring of incoming and/or outgoing calls, call redirection, and network user identification are selectable on a per subscriber basis. The flexibility of the NSNs are further enhanced by allowing all port parameters to be changed on-line, without affecting network operations or the other ports on the node. The NSNs are designed to solve your network problems.

Traditional Product - N7400

The Nucleus 7400 is the original member of the Nucleus 7000 Series. The N7400 incorporates a built-in diskette drive for fast, local software and configuration loading. Its real-time operating system (PPMOS) enables the workstation and LP cards to work in harmony, solving critical networking problems. Up to 32 high-speed ports are available in 4-port increments. Each port is software selectable and configurable on a port-by-port basis.

Figure 1-2 illustrates an N7400 Network Switch with a Standard LP card and a T-1 LP card. Various combinations of LP cards can be intermixed in the same workstation. Synchronous and trunk port speeds up to 256,000 bps and asynchronous speeds up to 38.4 Kbps are available using the Standard LP card. By using the T-1 or E-1 LP cards trunk speeds of up to 2.048 Mbps are possible. This performance is attained because all networking operations including control of the physical port interface, protocol execution, and packetized the data are performed by the LP card - freeing the workstation to perform its operations.

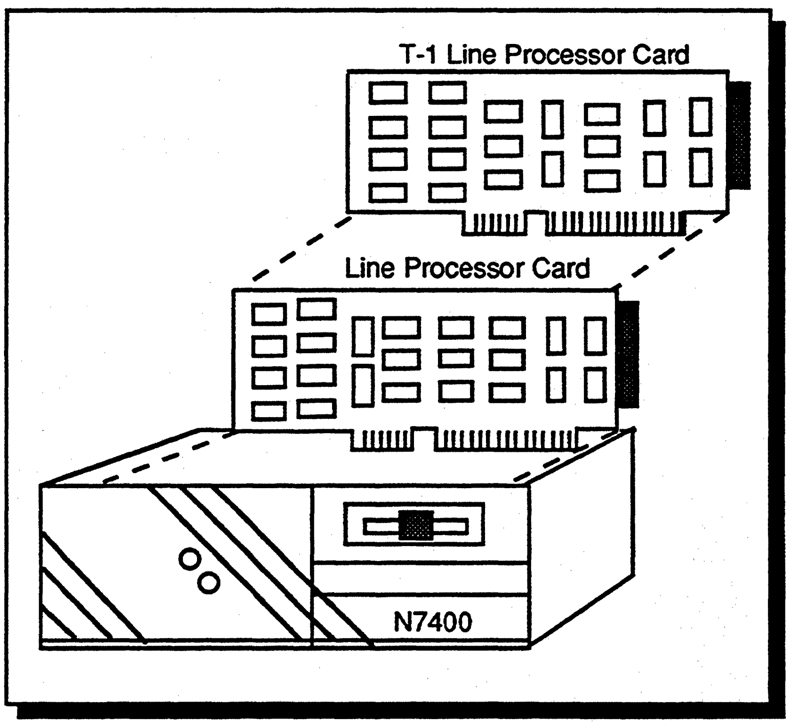

Figure 1-2 N7400 Configuration

The N7400 is designed to be adaptive to changing user and networking requirements. The N7400 supports a wide variety of protocols and can be configured to meet your specific access and throughput requirements. The following protocols are an integral part of the N7400: X.25, Async PAD (X.3, X.28, X.29), and SNA/SDLC PAD. The N7400 supports the MS-DOS environment, operates in 640 Kbytes of memory and supports 264 simultaneous connections.

The NSNs are adaptable to various devices through the assignment of a set of parameter on a port-by-port basis. Other parameters control the line protocol and data transferred between a terminal user and the computer. A dynamic adaptive routing algorithm based on Arpanet's routing algorithm, large packet sizes from 16 up to 4,096 bytes, automatic congestion control, and load balancing are just some of the features incorporated into the N7400.

The UNIX™ Connection - N7500

The Nucleus 7500 is the next generation of the Nucleus 7000 Series Network Switches. The N7500 uses UNIX™ as the base operating system for the workstation. This provides you with the ability to support larger networks; supporting up to 2,000 simultaneous connections per node and support for user applications like Integrated Call Processing (ICP). ICP minimizes call verification and call setup times and is designed for large networks, diverse networks, or where call setup times are critical. This optional software package is only available on the N7500.

Integrated Call Processing allows the N7500 to support call processing and call setup functions locally, without involving the N7900 Network Manager. It can also serve as a call processor for N7400s, that are assigned to it, so that they need not access the N7900. ICP contains the entire network database and participates with the N7900s in the database concurrency operations. Figure 1-3 shows how ICP can help improve overall network performance and integrity. The number of hops is dramatically reduced as well as the time to establish a call.

On a typical network the user would place a call for an application resident on the host computer. The call would follow the route outlined by the number one (1) path in order to verify the subscriber ID. Once the subscriber is verified, the call can be placed via the best route. Through the ICP application task the same user could use the capabilities of the N7500 for the verification process via the path outlined by number two (2). After the call is verified, the call is established via the best route. Imagine the power of this on large and diverse networks. Overall network performance would be improved by incorporating several N7500s throughout your existing network. The N7500, like the N7400, supports large packets up to 4,096 bytes for any network application or user requirement.

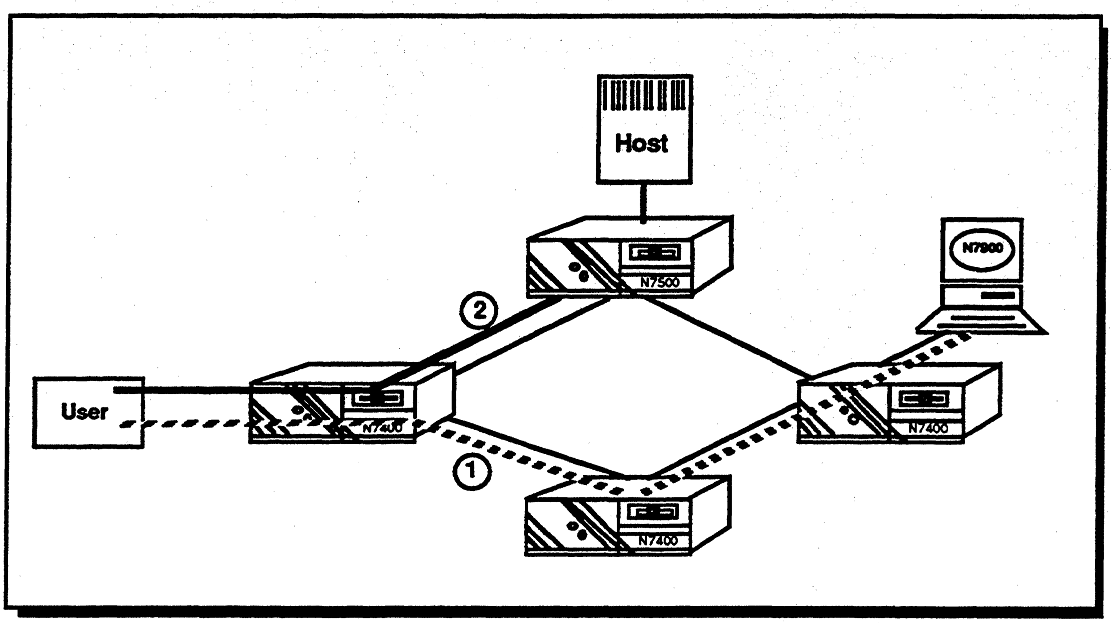

Figure 1-3 Integrated Call Processing

The N7500 includes a disk subsystem with a diskette drive, an 80 MB hard disk for configuration loading, call accounting information, and network database topology. The N7500 also supports a variety of LP cards allowing you to tailor your network for specific applications. Our Standard LP card supports trunk and synchronous operations up to 256 Kbps and asynchronous operations up to 38.4 Kbps. Standard line interfaces supported are RS-232-C (V.24/V.28) or optional RS-530 or V.35.

Higher throughput is obtained by using our T-1 or E-1 LP cards. Support for fractional T-1 lines are also available. Each T-1/E-1 LP provides up to four logical ports that operate over the T-1 or E-1 service. T-1/E-1 LPs support the *Fast Trunk* and *Standard Trunk* protocols. The maximum speed of each logical port is 1.544 Mbps for T-1 lines and 2.048 Mbps for E-1 lines, respectively. Additional logical channels can be added to the initial configuration through the use our 4-Channel Expander card. Figure 1-4 shows a typical N7500 configuration.

The N7500 includes state-of-the-art dynamic and adaptive data routing as well as automatic load balancing on parallel trunks. Congestion and avoidance control are inherent, ensuring a high degree of reliability and cost-effective data transmission across the network. To enhance transmission performance and back-up capabilities AMNET provides three (3) additional internodal trunk protocols besides to our standard LAPB protocol. These protocols are:

- Fast Internodal Trunk Protocol (Fast Trunk)
- Public Data Network (PDN) Trunk Protocol
- Back-up Trunk Protocol

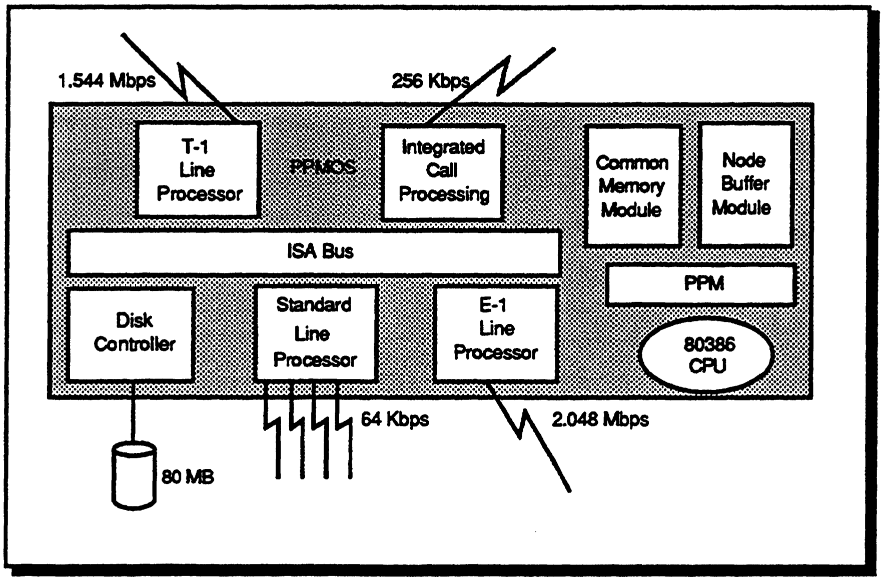

Figure 1-4 N7500 Configuration

These protocols are software selectable on a port-by-port basis. They can support analog lines as well as digital transmission facilities. When digital facilities are available we complement their inherent features of high noise immunity and error free transmission with our *Fast Trunk*. AMNET's *Fast Trunk* is available on both NSNs and can be used as an alternative to our standard LAPB trunk protocol. The *Fast Trunk* protocol improves throughput by 50 to 75 percent, reduces network delay, and dramatically improves user response time.

The Network Manager - N7900

The Nucleus 7900 (N7900) Network Manager is a companion member to either of the NSNs. The N7900 is an important part of AMNET's *Integrated Network Solution* and provides all the management functions critical to running the data network. Its color graphics user interface is utilized in all modes of operation including configuration management, monitoring, administration, and database management. It is connected to an adjacent network switch via a proprietary high-speed Parallel Interface Adapter (PIA). The PIA constantly monitors operations between the N7900 and its adjacent node whether it be an N7400 or N7500.

The Nucleus 7900 is based on the powerful UNIX™ operating system. The hardware configuration consists of the workstation, color monitor, 8 Mbytes of RAM, a 80 MB hard disk, floppy disk drive, a PIA module, and a serial printer. The N7900 supports the network configuration database, call verification, call setup and call accounting information, and continuously collects data on the status of every device and transmission circuit in the network. Figure 1-5 represents a typical network managed by the N7900.

The N7900 monitors the status of the nodes in the network and provides immediate visual indication of the location and severity of alarms, allowing operators to respond rapidly to changing network conditions. Network operators are alerted to alarms by changes in color of the component (icons which represent subnetworks, N7400 or N7500 Network Switches, or another distributed N7900s), as well as trunk lines. Furthermore, a hard copy of all alarms and events is generated on the printer attached to the N7900 as well as being archived to the hard disk. Alarm status is defined according to thresholds set by the network administrator.

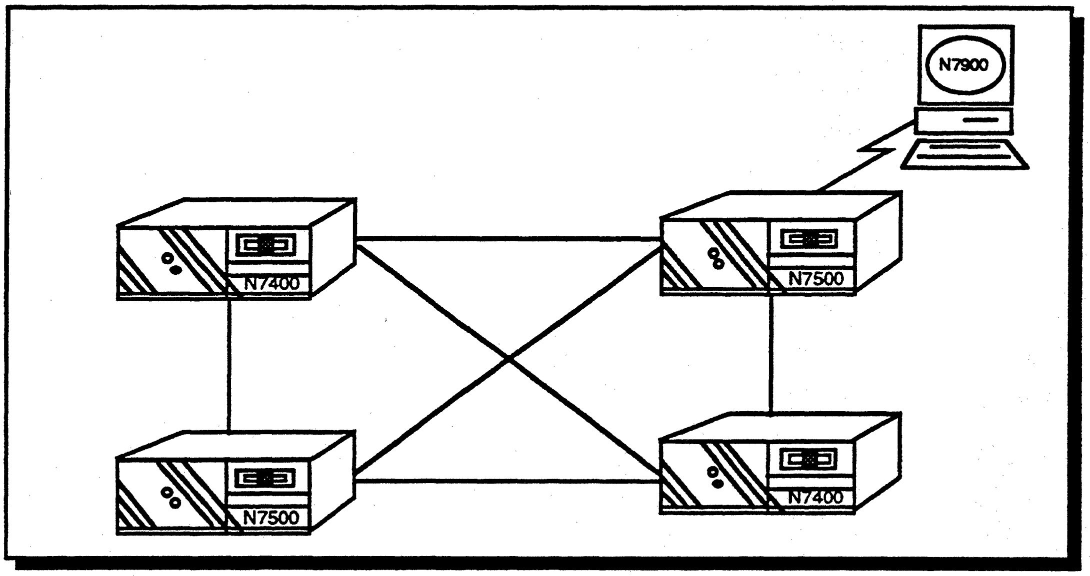

Figure 1-5 Typical Backbone Network

All events are time stamped and automatically archived to the hard disk. Now operators can use this information to predict potential problem areas, spot problem conditions, and produce historical reports. The network operator can zoom-in to a faulty node or even down to the physical port level, visually inspect the EIA leads and diagnose the problem. Once the problem is diagnosed remedial action can be taken to correct the problem.

In summary the N7900 Network Manager provides:

- Network configuration management
- Real-time statistics gathering and analysis
- Automatic capture of all network events and alarms
- Support for stand-alone or redundant network managers
- Powerful Graphics to visualize network topology and pinpoint errors

The advanced system architecture of the N7900 provides you with the ability to control all network resources from a centralized location or control can be distributed throughout your organization at regional or campus locations. For diverse networks, that span time zones or countries, all alarms and events can be automatically forwarded to other N7900s located throughout your network. This allows local network operators to view the problem and take corrective action.

At any point in time there is only one network master with a master database controlled by the current N7900. If modifications are made to the master database, these modifications are recorded on the N7900. When several N7900s or N7500s are located throughout the network, their individual databases are automatically updated to reflect the changes made to the master database; ensuring database concurrency.

X.25 Facilities

The NSNs provide an integral X.25 communication interface and frame Level III switching capability in accordance with the 1984 CCITT Recommendations. The Release 3.0 enhancements encompass all the essential 1984 X.25 facilities. The essential facilities are realized by moving some of the per subscriber and network facilities to the individual port level. Any port or multiple ports can be software configured as an X.25 port and each port supports a wide range of characteristics that can be optimized to increase its performance. Both Switched Virtual Circuit (SVC) and Permanent Virtual Circuit (PVC) services are supported. The following port enhancements include:

- Lowest Incoming Channel
- Highest Incoming Channel
- Lowest Two-way Channel
- Highest Two-way Channel
- Lowest Outgoing Channel
- Highest Outgoing Channel
- Default Packet Window
- Default Packet Size
- Maximum Packet Size
- Flow Control Negotiation
- Throughput Class Negotiation
- RPOA Selection
- Default Calling Address
- Local Charging Prevention
- Transmit Delay Selection and Indication

On a per subscriber basis the following facilities are included in Release 3.0:

- Network User Identification (NUI)
- Calls Barred - Incoming and Outgoing
- Call Redirection
- Fast Select Acceptance
- Reverse Charging Acceptance
- Hunt Group
- Closed User Group

This level of flexibility provides an important tool to network designers in optimizing system and network performance. The process of integrating the network with other X.25 devices is also simplified by having control of key interface characteristics.

Multiple Trunk Protocols

AMNET provides you with the ability to customize your network as well as your access into the network. In dealing with diverse networks you can select any of the following trunk protocols to support your networking requirements:

- Standard Trunk
- Fast Trunk
- PDN Trunk
- PDN Backup Trunk
- Backup Trunk

The *Standard Trunk* supports the Standard LP cards from 9.6 Kbps to 256 Kbps. This is the default trunk protocol for any using our Standard LP card. It uses the CCITT X.25 LAPB procedures to support access over analog or questionable transmission lines.

The *Fast Trunk* protocol is an OSI Layer II protocol. It uses a subset of the full CCITT X.25 LAPB procedures used by our *Standard Trunk* protocol. This protocol is based on the frame relay technology, that is, this protocol performs error detection only with no error correction and is utilized over reliable digital facilities. Any packets received with errors are simply discarded. Packets with errors will be retransmitted by the source node. Like all the other protocols offered on the NSNs, the *Fast Trunk* protocol is also software selectable.

By selecting the *Fast Trunk* protocol, the network operator may see a 50 % to 75 % improvement in throughput over the *Standard Trunk* protocol. T-1 and E-1 facilities are typically very reliable and error free. Therefore, the default internodal protocol when using any T-1 or E-1 LP card is the *Fast Trunk* protocol. This can also be changed by the network operator to use the *Standard Trunk* protocol.

PDN Trunk support provides for the consolidation of a multinational, non-contiguous network where end-to-end private service is not available or is too costly to implement. You can install nodes and private lines where it's cost-effective; that is, you may install a private network domestically and another private network at your international office in Tokyo. But the only feasible way to interconnect these two networks is through multiple PDNs. AMNET's End-to-End protocol traverses PDNs to ensure data integrity and delivery of the information on an End-to-End basis. Figure 1-6 illustrates such a network incorporating NSNs and our N7900.

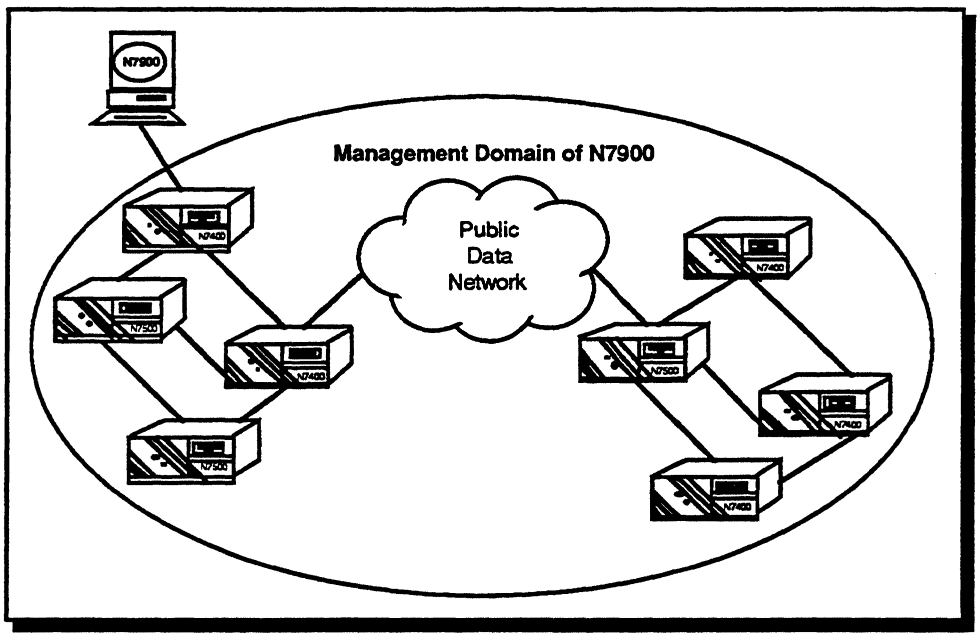

Figure 1-6 Management Domain

Another feature of the *PDN Trunk* is its gateway facility. As a gateway port, it can translate portions of the Call Request packets based on the contents of the packet and operator entered translation formulas. Each port has two translation tables associated with it; one for inbound Call Request packets and one for outbound Call Request packets. Translation templates consist of a Template ID and a set of operations to apply to the called address, calling address, user data field, and the facility fields of Call Request packets. Each of these fields can be passed unchanged, deleted entirely or replaced with another value. In addition, 'formula' operations can be applied to the called address and user data field.

Since both the N7400s and N7500s provide dynamic routing of packets around faulty or congested nodes, the capabilities of using *PDN Backup Trunk* is not required. But in the event of a catastrophic error, in which dedicated trunk lines experience a failure, you are provided with the added protection of a back-up path, namely through the PDN facility. Autodial modems used in conjunction with the NSNs will automatically reconnect adjacent networking nodes using the PDN as the transport network. This is a cost-effective means of implementing disaster recovery and fault tolerant operations.

The power and capabilities of the NSNs are further enhanced through the use of the *Backup Trunk*. In fact, the network operator can establish up to eight logical or *Backup Trunks* per trunk line. The *Backup Trunk* supports the Standard LP card as well as the T-1/E-1 LP cards. In the event of an error on a trunk line, your node will automatically use a back-up path directly through the NSN as illustrated in Figure 1-7. All calls will stay active during the switch over to the *Backup Trunk*.

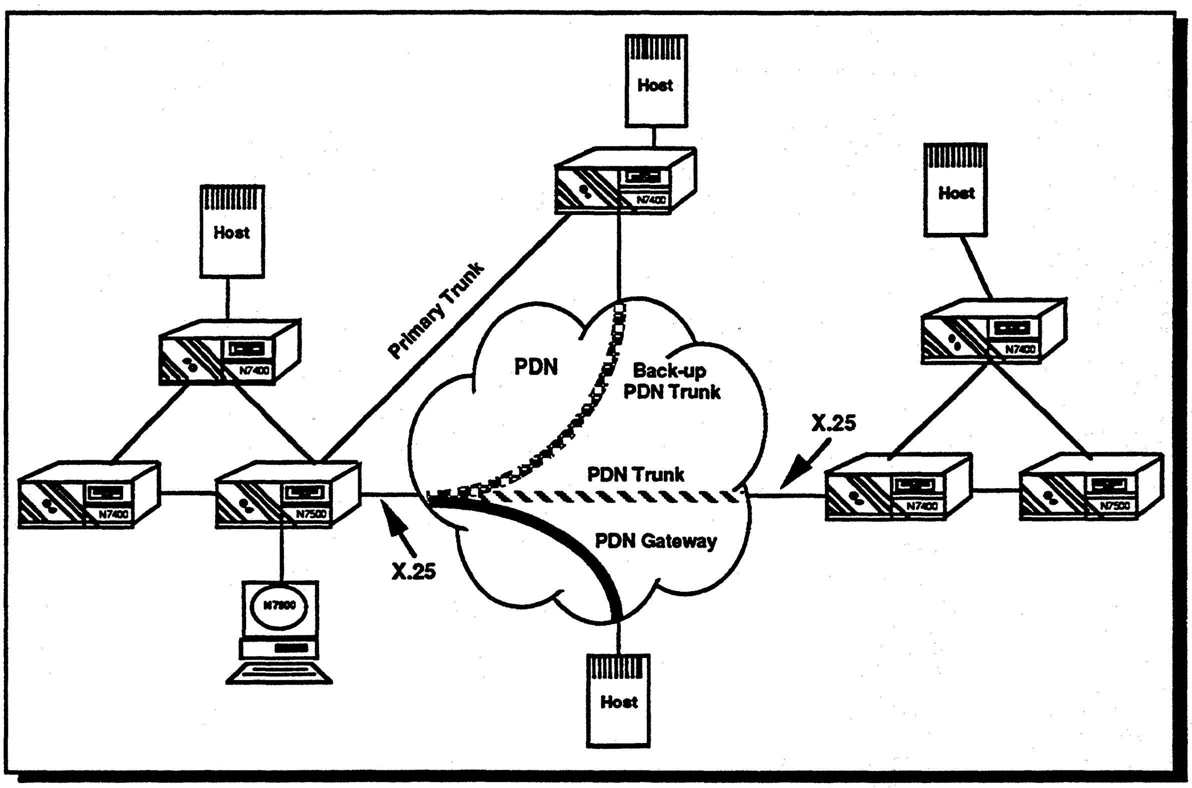

Figure 1-7 PDN Backup Support

AMNET's End-to-End Protocol

AMNET's *End-to-End* protocol provides for guaranteed delivery of information between the source and destination nodes in the network. This protocol is a transport level protocol similar in functionality to the ISO Transport Protocol Class 4 (TP4) and DOD TCP protocols. To ensure data integrity and reliable transport of information, AMNET's *End-to-End* protocol automatically performs resequencing of out of sequence packets, retransmission of missing packets, detection of duplicate packets, and *End-to-End* flow control. Your applications and data are further protected with additional error correction and control facilities. Figure 1-8 shows where in the ISO structure AMNET's *End-to-End* protocol resides.

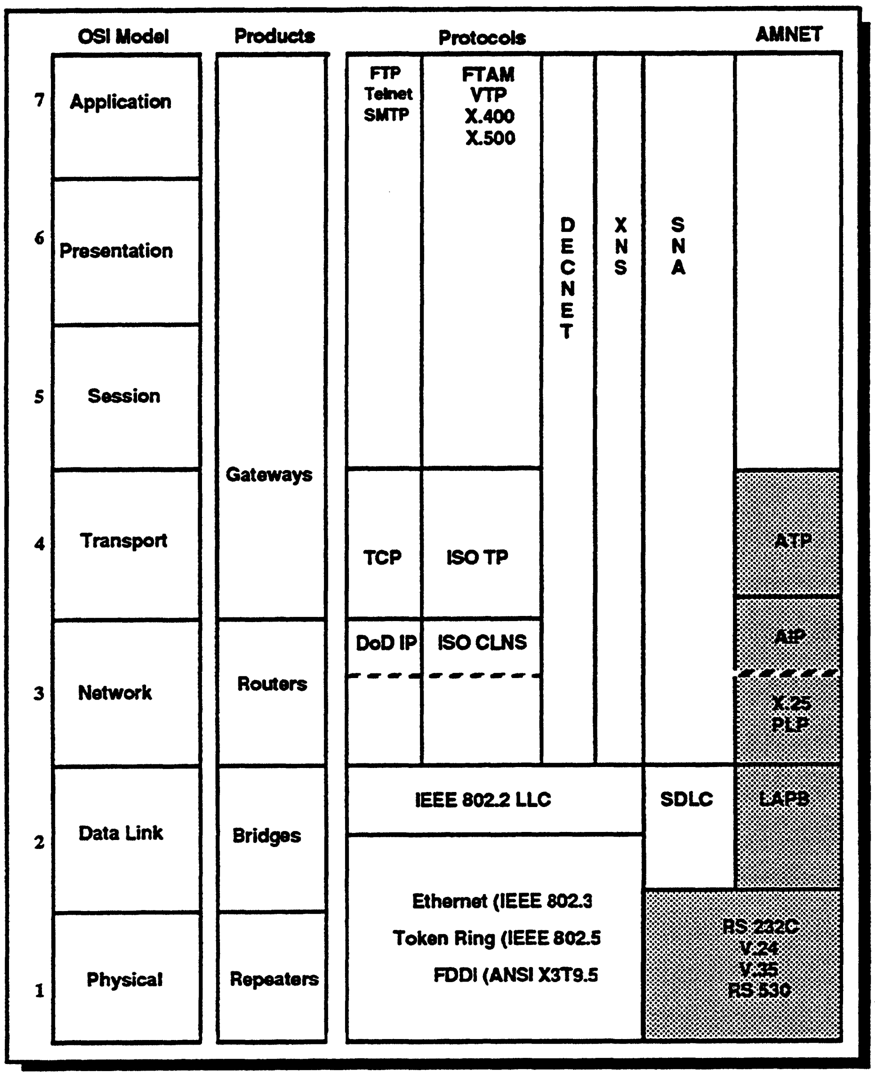

Figure 1-8 AMNET's End-to-End Protocol

Line Monitoring Mode

A unique capability standard of AMNET's NSNs allows the operator to monitor any incoming or outgoing line on the node in a data scope fashion. Any N7400 or N7500 that has a keyboard and display monitor can use this feature as a local data scope. With the *Line Monitoring Mode* you can display all packet level events received and transmitted by the node. Packets received by the node are in normal video and transmitted packets are displayed in reverse video.

To isolate the problem you can specify an LP and port to be selected for monitoring, or a Group Send message can be monitored as well as test messages generated by the N7900. The data can be viewed in ASCII or in hexadecimal formats and the operator can selectively toggle between both formats. The operator can further isolate the LP and port by selecting one LCN, on that port, to monitor.

Trunk lines display the trunk packet protocol and all other lines display the X.25 packet protocol. For PAD ports the display shows the X.25 output from the PAD port.

The Asynchronous PAD

An integral Asynchronous PAD facility provides transparent communication between remote or locally attached terminals, printers, and workstations to one or more central computers over the NSN. The Async PAD conforms to CCITT 1984 3X standards (i.e., X.3, X.28, and X.29).

Any port can be software selected and configured as an Async PAD port. Unlike other Async PADs, we allow you to select port speeds up to 38.4 Kbps. The Async capabilities are further expanded to include support of our 4-port Daughter card. Additional features include:

- RTS/CTS Flow Control
- Welcome Banner Per Port
- Auto Disconnect with Configurable Timer
  - If no login
  - If no command (after login)
  - If no data (after command)
  - If clear received by the PAD
  - If clear sent by the PAD
- Drop Modem Signals on Disconnect
  - Configurable hang up duration

The SNA PAD

The N7400 and N7500 also supports IBM's Systems Network Architecture (SNA) via an integral SNA/SDLC PAD. This facilitates communications between remote SNA Control Units and an SNA host processor over the backbone network. Any port on the NSN can be configured as a SNA PAD. Configurable PAD types include the TPAD and HPAD. The logical channel connection between an HPAD and TPAD can be either a PVC or a SVC.

The SNA/SDLC Terminal PAD, or SNA TPAD, acts as a gateway for SNA users over the NSNs. It provides communications between remote SNA Control Units and an IBM host or compatible computer connected to the X.25 network through an SNA HPAD or NPSI software. The TPAD supports Qualified Logical Link Control (QLLC) procedures.

The SNA/SDLC Host PAD, or SNA HPAD, acts as a gateway for SNA host computers and an X.25 network regardless of whether it is a private or a PDN. It offers communications between an SNA host and remote SNA devices having X.25/QLLC support or connected to the X.25 network through an SNA TPAD.

Both the TPAD and HPAD perform automatic call setup and packetizing of SNA information frames. The SNA TPAD performs local polling of the Control Unit and the HPAD responds to local polling from the host, thus keeping network traffic to a minimum. The HPAD supports up to 16 SNA TPAD ports connected anywhere on the X.25 network and each TPAD port can support one Control Unit. Hunt Group facilities can be configured for a set of HPADs connected to the same IBM or compatible Front End Processor (FEP) providing alternate routing.

AMNET's additional functionality provides support for NPAD and XPAD. NPAD provides IBM Host-to-Host support. XPAD performs transparent PAD operations, that is, it provides the pipeline through which two SNA nodes; using the Q-bit of the X.25 data packet; can transfer SNA control data. The QLLC provides the services of the link layer as perceived by higher levels of SNA architecture. For example, the connection between a PU Type 2 node (SNA 3270/SNA 3770) and a PU Type 4 node with/without supporting the X.25 functions (37x5 with/without NPSI under the NCP). Both XPAD and NPAD requires a one-to-one relationship.

The Elements of the NSN Family

Each NSN can include the following: one (1) to eight (8) LP cards either the Standard 4-port LP card or our T-1/E-1 LP, an optional 4-port Async only Daughter card which mounts onto the Standard LP card, industry standard 386-based workstations with an ISA Bus architecture, and AMNET's designed and developed real-time operating systems, tasks, and applications that control and manage the LP card as well as the resources of the workstation. Figure 1-9 illustrates the hardware and software complement of AMNET's Network Switches and Network Manager.

An optional 19-inch rack mount enclosure is available for all NSNs, Hot Standby Switch, as well as the N7900. When the rack mount enclosure is used for the NSN it provides additional expansion up to a total of 64-ports. It can also be used in hostile environments with its larger power supply and superior cooling requirements or where a 48 volt DC power supply is required. The 48 VDC power is mandatory in many government facilities, telephone central offices, and large corporations where network access and uptime are paramount.

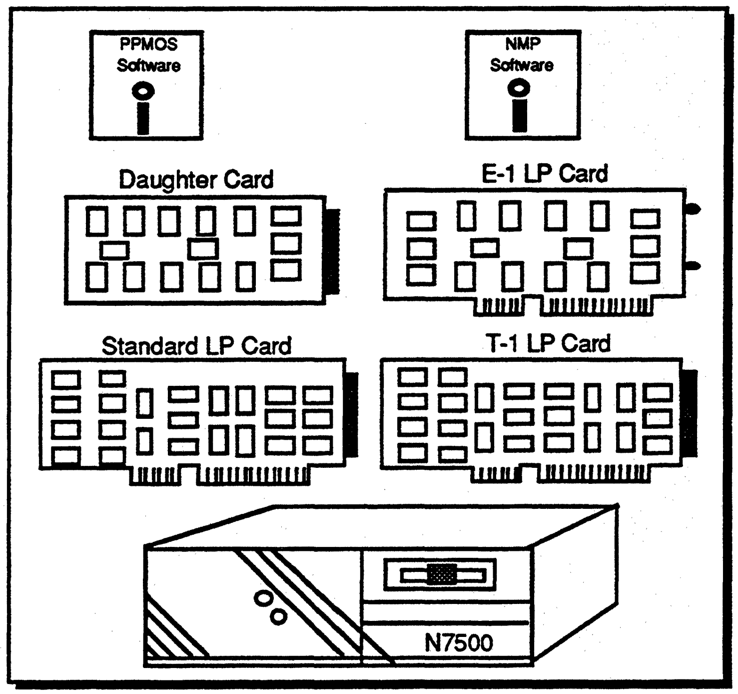

Figure 1-9 NSNs Major Components

All AMNET developed applications and real-time operating systems are written in the 'C' language for portability and ease of maintenance. There are two major software components of the NSNs.

First, a real-time operating system which manages the LP cards and supplies the interface to the communication ports. All LP applications run under the Front End Processor Operating System (FEPOS). This is an AMNET developed, interrupt driven, real-time operating system providing the task dispatching, memory management, and I/O and interrupt handling for all the LP cards located in the node.

The FEPOS software is responsible for the physical (Level I) and logical and link (Level II) interfaces. Further it contains the LAPB (trunk, *Fast Trunk*, and X.25), the 3X PAD (X.3, X.28, and X.29), and the SNA/SDLC PAD applications. The operating system and applications are resident on each LP card providing total independence and autonomy. During workstation initialization each LP card is downloaded with its own copy of the operating system and designated applications. Therefore, should any one LP card fail only the subscribers on that LP would be affected and not the rest of the network.

Second, the Packet Processor Module (PPM) software operates in the workstation and provides the packet switching and user interface capabilities. All of the application tasks running in the packet processor environment run under the PPM Operating System (PPMOS). This operating system provides task dispatching and memory management for the workstation environment. Optimal networking performance is obtained by having all common data structures accessed through operating system calls and communication with the LP cards are handled through a shared memory region. (Shared memory is also referred to as dual-ported or multi-ported memory.)

NSNs Characteristics and Features

The N7400 and N7500 Network Switches provide the following:

General Characteristics

- Up to four (4) software configurable ports per LP card
- Maximum aggregate throughput of 256,000 bps per Standard LP card
- Maximum packet size supports up to 4,096 bytes
- Sustained throughput rate of 1,200 packets per second (pps)
- Supports Switched Virtual Circuits and Permanent Virtual Circuits
    - Supports up to 264 simultaneous VCs on N7400
    - Supports up to 2,000 simultaneous VCs on N7500
- Supports CCITT X.121 addressing
- Connects 25 - 40 calls per second (Configuration dependent)
- RS-232-C (CCITT V.24/V.28) interface
- Operating configurations loaded upon power-up and stored on disk

X.25 Interface

- Compatible with 1984 CCITT Recommendations for X.25
- Configurable as DTE or DCE
- Line speeds up to 2 Mbps on each synchronous or trunk port
- X.25 facilities supported:
    - Closed User Group
    - Fast Select
    - Reverse Charging

- Default Calling Address
- Local Charging Prevention
- Window and Packet Size Negotiation
- Flow Control and Throughput Class Negotiation
- RPOA Transit Network Selection
- Network User Identification

• Extensive Call Control capabilities:

- Hunt Group
- Group Send Facility
- Call Redirection
- Priority Request
- Call Barred (Incoming and Outgoing Calls)

Asynchronous Interface

• Baud Rates from 50 to 38,400 bps
• Local Login Password for PAD users
• Autobaud (1,200 to 19,200 baud) and autoparity
• Dedicated line, dial-in, and dial-out connections
• Accommodates CCITT X.3, X.28, and X.29 parameters
• Downloadable Profiles
• Welcome Banner on a per port basis
• RTS/CTS hardware flow control
• Abbreviated Addressing
• Auto Disconnect with configurable timer

- If no login
- If no command after login
- If no data after command
- If clear received
- If clear sent

SNA/SDLC Interface

• Up to 64 Kbps per port for frames up to 280 bytes
• SNA/SDLC ports can operate as DTE or DCE
• Data link control: SDLC full-duplex mode at Level I and half-duplex at Level II

- Configurable SDLC frame size
- Maximum frame size is 4,105 bytes
- Configurable station addresses
- Point-to-Point operation for TPAD connections
- Support of NRZ or NRZI
- Autocall on both HPAD and TPAD
- QLLC/NPSI compatible for HPAD and TPAD connections
- Supporting four modes:
  - HPAD
  - TPAD
  - NPAD (Host-to-Host)
  - XPAD (Transparent)

Nucleus 7400 and Nucleus 7500 Features

The features of the N7400 and N7500 are as follows:

- High performance to 1,200 pps
- Dynamic Adaptive Routing
- Maximum Async port speed up to 38.4 Kbps
- Maximum synchronous or trunk port speed up to 256 Kbps
- Maximum internodal trunk speed up to 2.048 Mbps
- *Fast Trunk* for digital facilities
- Integrated Call Processing (N7500 only)
- Automatic Load Balancing over parallel paths
- Dynamic non-disruptive routing around line congestion and failures
- Up to 64-ports per NSN
- Implements CCITT X.3 parameters (Xon/Xoff, RTS/CTS)
- Implements CCITT X.28 and X.29
- NSPI and QLLC compliant SNA PAD ports
- PDN Gateway support
- Menu-driven, user-friendly operation
- PDN Back-up support over public or private data network

Nucleus 7900 Network Manager Features

- Comprehensive centralized or decentralized monitoring and control
- Powerful Graphical User Interface to pinpoint errors
- Color coded alarms and events, up to four (4) levels depending on severity
- Color coded maps to visualize network topology
- Remote network configuration management
- Remote software distribution
- Real-time statistics gathering and analysis
- Automatic capture of all network events and alarms
- On-line problem detection and fault isolation
- Menu-driven support with on-line Help facility

Applications

The NSNs are used in numerous applications. On the following pages we show a sampling of some of the industries and applications where the NSNs are used to solve our customers requirements.

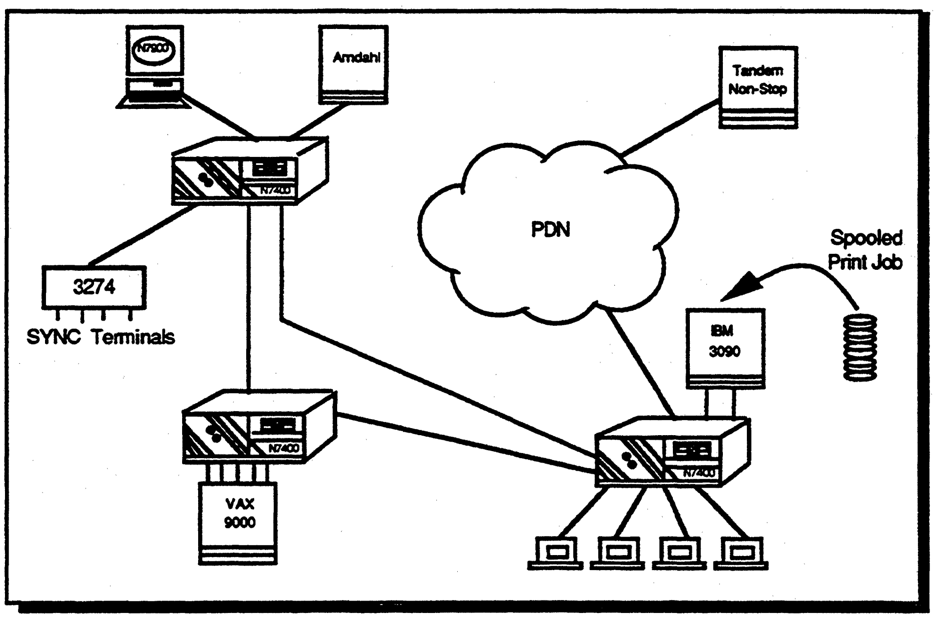

Figure 1-10 Heterogeneous Computer

Connecting Heterogeneous Computers

Bankers, lawyers, and manufacturers are experiencing problems of data overload and accessing the required information they demand from various sources. Government and state agencies have standardized on various computer vendors. The task of integrating these diverse systems and technologies into a homogeneous network is now required.briefly highlight a portion of these tasks to provide you with a better understanding of their capabilities and some of the capabilities of the N7900 Network Manager. (Programmers, system designers and system integrators that require greater detail can reference the *LP and PPM Software Programmer's Guide.*)

AMNET's N7400, with integrated PAD support, allows both asynchronous and synchronous devices to communicate across a common network allowing users access to these various systems. Diverse computer systems, workstations, remote terminals, or PCs can be connected together to support your user community. Furthermore, they can be located at local or at remote locations. Figure 1-10 shows a typical application connecting heterogeneous computers and devices.

Dedicated Backbone Networks

The flexibility and expandability of the NSNs are illustrated in Figure 1-11. Remote and local terminals, PCs, and workstations are connected via the integrated PADs allowing users to access applications on several hosts, over PDNs, or to each other. *Fast Trunk* protocol is implemented on digital facilities as well as T-1 LP cards for internodal networking requirements for high reliability and bandwidth.

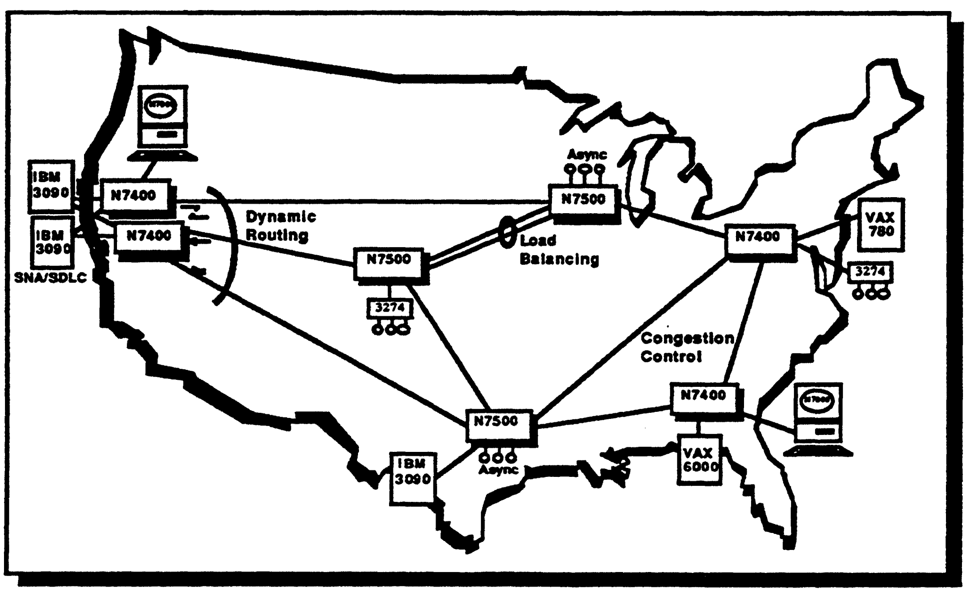

Figure 1-11 Backbone Networks

With multiple nodes, alternate paths are specified so if a trunk is unavailable or lost, traffic is automatically routed over an alternate or back-up path to its destination via the best path. For any subscriber defined by the Hunt Group up to eight (8) paths can be specified as alternate routes, with a choice of three (3) algorithms to optimize networking conditions.

Every NSN in the network can be controlled from the N7900 located at the central site. To facilitate network management, all events and alarms are directed to the N7900 and printed on the serial printer. The PVC capabilities of the NSNs can be used to automatically establish routing and connections between the N7400s and N7500s.

Private Networks

International companies have established private networks for various reasons: namely; security, availability, and expansibility. Figure 1-12 shows an international company that uses some of the capabilities of the N7400. The criteria for the network that had to be met were as follows: first, 99% uptime and availability; second, a Network Manager to handle alarm and statistics gathering; and third, the ability of the network to easily be reconfigured.

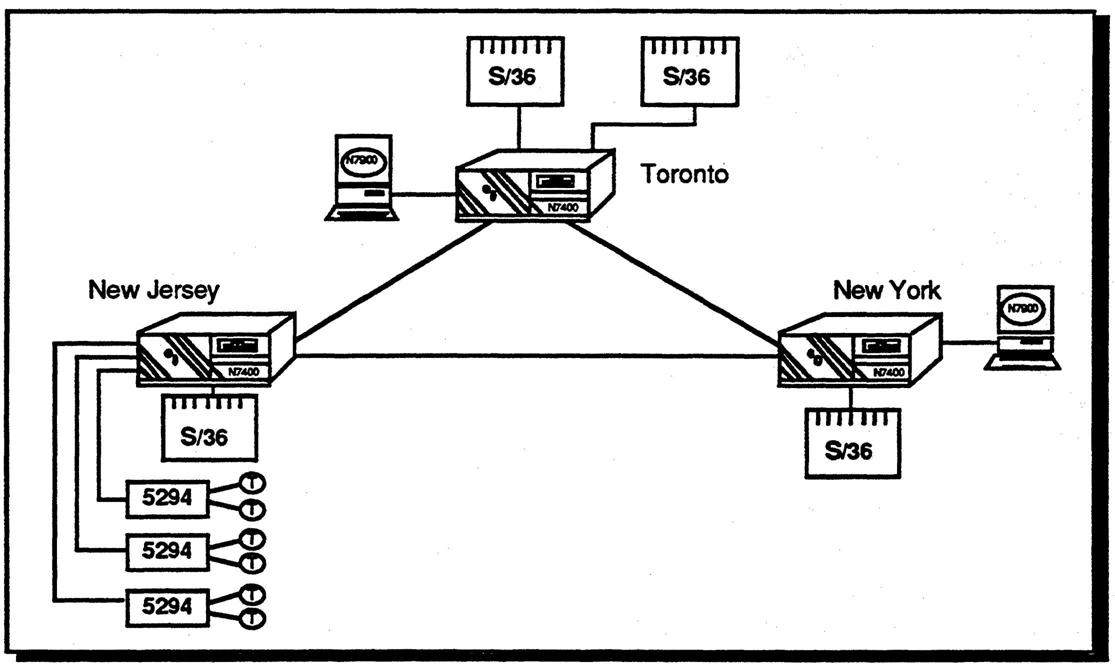

Figure 1-12 Dedicated International Network

The network is used by both the marketing and financial departments in a peer-to-peer communications applications for file and data exchange and in remote control applications for printing. Marketing uses the system with some on-line applications which produces product information which is distributed to the sales force. Finance applications included the upkeep of general ledgers and sub-ledgers; accounts receivable; invoicing; and order acknowledgements. The benefits derived from using the N7400s are the ability to keep the user application up and running even with a line failure; provide an upwards migration path to support more users; and the ability to have distributed control centers for the network and its resources.

Public Data Networks

A majority of European companies have standardized on using PDNs as their backbone network. One of the primary reasons behind this move is the premium cost associated with dedicated high-speed lines. Another reason is the changing requirements and duty imposed once the data crosses demographic boundaries.

AMNET has enhanced the capabilities in our NSNs in this area as well as provides an automatic back-up facility incorporating the PDN. Figure 1-13 encapsulates some of the capabilities of the N7400 and N7500 in supporting and utilizing the PDN facility.

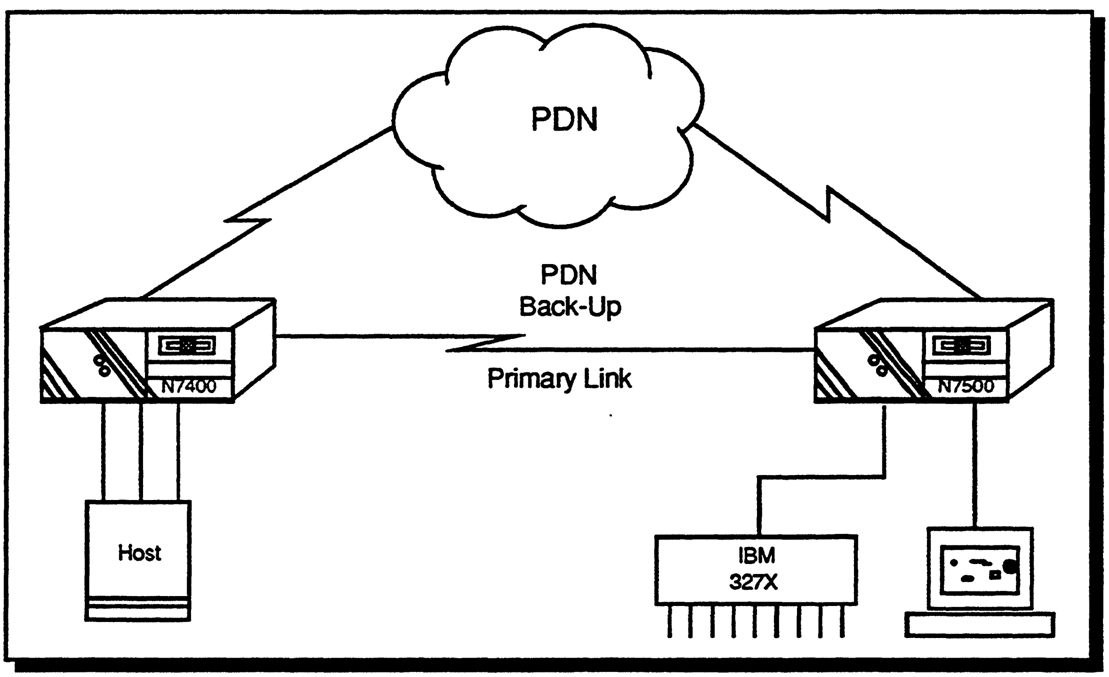

Figure 1-13 Public Data Networks

SECTION 2

ARCHITECTURE

&

THEORY OF OPERATION

Node Architecture & Theory Of Operation

In this section, concepts, terminology, and architecture are described in detail to illustrate the relationships that exist in the following operational areas:

- Node Hardware and Software Architecture
- Line Processor Cards
- Asynchronous, and SNA/SDLC Device Support
- N7400 & N7500 Windows and Flow Control

The NSN Hardware Architecture

Both the N7400 and N7500 are high performance Network Switches. Their high performance is obtained by adherence to software design goals and a modular hardware design. The hardware architecture is based on a powerful 80386-based workstation coupled with intelligent LP cards. The N7400 is incorporated on the 386 workstation, with 1 Mbyte of memory, and a floppy disk drive. The diskette holds the configuration parameters for the node and automatically loads these parameters on the node when the node is brought on-line or powered up.

The N7500 also uses a 386 workstation but with 8 Mbytes of memory, an 80 MB hard disk and a floppy disk drive. The diskette drive is used to load application programs and as a back-up facility. The hard disk has the UNIX and AMNET operating systems, configuration parameters for the node, and additional user applications. When the node is powered on or rebooted the LP cards are downloaded with the nodal configuration from the hard disk. With the additional memory, larger buffer space is available so that the N7500 can support 2,000 simultaneous virtual circuits. Also, user application programs like the Integrated Call Processing are available.

Our family of high performance LP cards are intelligent and supported by the NSN. First, our Standard LP card (N7400-E04) has 4-ports with an aggregate throughput of 256 Kbps and can be configured for Async, trunk or synchronous operations. Each port is software selectable on a port-by-port basis. Our modified LP card (N7400-S04) is identical to our Standard LP card but supports additional line interfaces; namely, RS-530 and V.35. To complement the Standard LP and modified LP cards is an Async only Daughter card.

This daughter card physically attaches to the LP card and cycle steals from the processor on the LP card. The daughter card does not require a ISA bus slot because it mounts on the LP card. However, it does require access to a back panel slot on the workstation to attach the 4-port splitter cable. Therefore, each LP card can support four (4) synchronous, asynchronous or trunk ports plus four (4) additional asynchronous only ports for a maximum of 64-ports for an eight (8) LP configuration.

To provide higher internodal bandwidth and multiple logical internodal trunks we have available LP cards that operate at 1.544 Mbps to 2.048 Mbps, respectively. These LP cards are our T-1 and E-1 Line Processor cards. The T-1 card supports line rate up to 1.544 Mbps, four (4) Nx56 or Nx64 Kbps logical trunk lines, and supports fractional T-1 (FT-1) operations from 56 Kbps to 1.536 Mbps. Two (2) models are available for your network. The N7450-001 supports an integrated Channel Service Unit (DSU/CSU) interface for direct access to any T-1 facility. The other model (N7450-000) has a DSX-1 interface and requires an external multiplexer and CSU.

For international operations, its counterpart is an E-1 LP card (N7450-002) which supports the G.703 standard and adheres to the CEPT recommendations. The E-1 LP card supports line rates up to 2.048 Mbps and four (4) Nx56 or Nx64 Kbps logical trunk lines where 'N' is from 1 to 30 or from 64 Kbps to 1.920 Mbps. A key feature with our family of NSNs and associated LPs is that you can freely intermix any combination of LP cards within our workstations. For example, the Standard LP cards can be used with our T-1 or E-1 cards in any combination up to a total of eight (8) LP cards. Figure 2-1 illustrates the capabilities of the network connectivity.

Both the T-1 and E-1 LP cards are expandable to support additional trunk lines in multiples of four (4) port lines. This is provided with our 4-Channel Expander Card. The Expander Card works in conjunction with a primary T-1/E-1 LP card. Also available is a two (2) port Sync/Serial Card. This card provides two functions: first, data can be switched from an incoming T-1/E-1 line onto this 2-port Sync/Serial card; and second, data terminal equipment (DTE) can be connected directly to a T-1/FT-1 or E-1 transmission facility. Some uses for this Sync/Serial card are for video conferencing and imaging equipment.

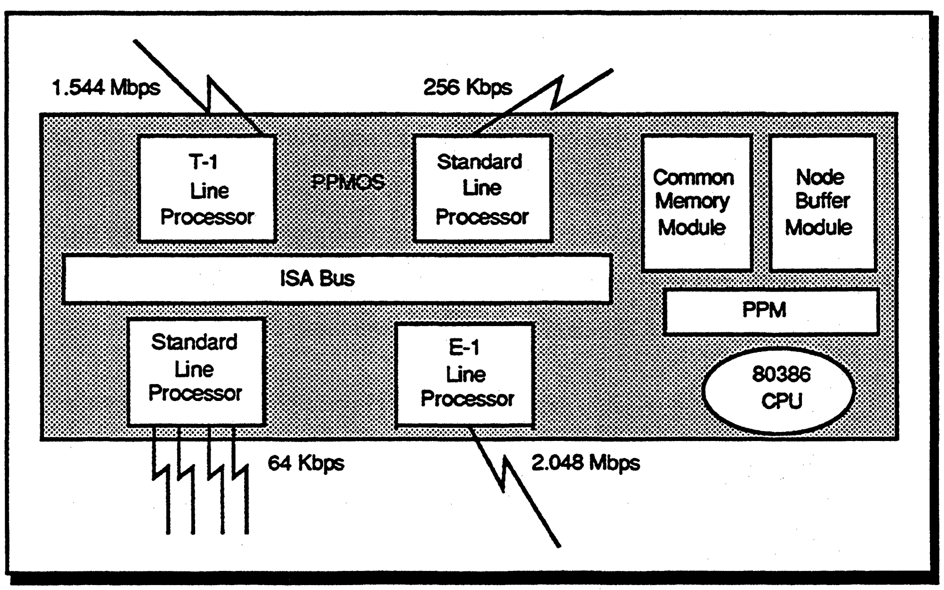

Figure 2-1 Network Connectivity

Intelligent Line Processor Cards

The intelligence of our LP cards are software developed and designed by AMNET. No matter which card you select for your network requirements the following hardware complement is included on each LP card:

- Powerful and dedicated 16-bit processor (80186 based cpu)
- 512 Kbytes on-board dynamic RAM memory per card
- 16 Kbytes of shared memory or multi-ported memory per card
- DMA I/O channels on all transmit and receive ports

The processor and RAM memory allows each LP card to operate in an autonomous fashion. This provides us with the ability to download each LP card with an individual copy of AMNET's operating system and application programs directly onto the LP card. Should one LP card experience a problem it does not effect the LPs or ports on the node. Performance enhancements come from the fact that we designed and developed the LP operating system (FEPOS) and applications that run on these cards. In fact, it is a simple and straightforward matter to develop other unique applications to run in the LP cards.

Dedicated data and buffer areas exist within each LP card to provide a common interface into the workstation. Table 2-1 lists these areas in the LP card and their size in bytes. These sizes are approximations since they will vary depending on the software release.

Name of Area	Size in Bytes
LP Programs, Data Area & Stack Space	286,000
LP Dynamic Buffer Pool	210,000
Shared Memory Region	16,000

Table 2-1 Line Processor Card Data Areas

By incorporating a shared memory scheme the sharing of messages, status information, and data are readily available with no performance degradation. The shared memory (multi-ported memory) region is constantly being updated with command and status information. Each LP and Packet Processor (PPM) poll the shared memory region to determine command and status information. The workstation and LP cards freely communicate by placing data, and status information into what they consider to be local memory.

Line Processor Card Characteristics and Operation

All Network Switching Nodes utilize the LP cards for all external I/O operations including back-up trunk, network trunk, subscriber access, and high-speed internodal trunk lines. Up to eight (8) LP cards can be used in the workstation including any combination of Standard LPs, T-1 LPs, and E-1 LPs. Figure 2-2 illustrates two LP cards and how the data flows through the node.

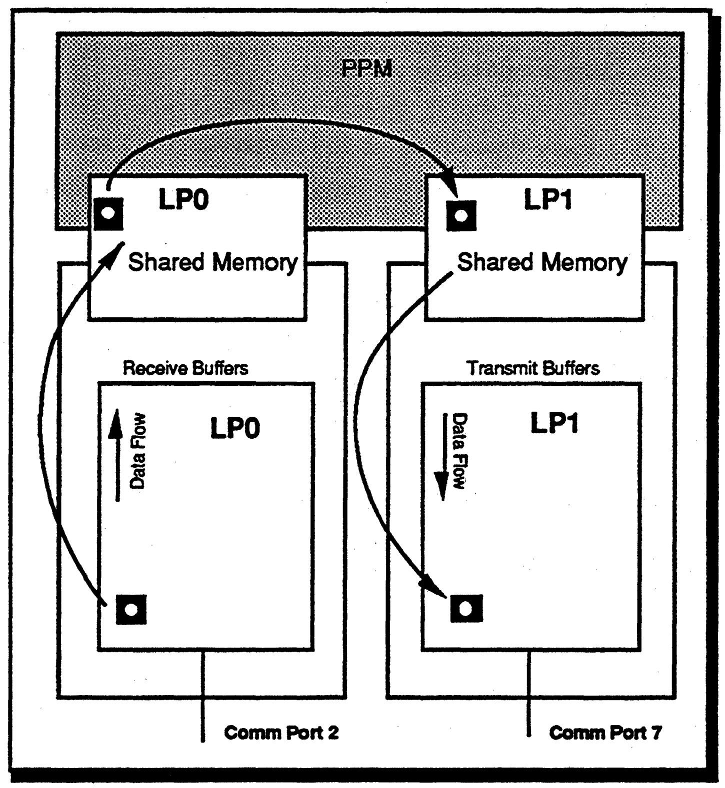

Figure 2-2 Typical LP Data Flow

Each LP has a dedicated area to hold the transmit and receive buffers and a shared memory area for communications with the workstation. The PPMOS runs in the workstation and provides the packet switching and user interface capabilities. All communications between the LP and PPM proceeds through this shared memory mechanism.

Inside the LP are dynamic buffers (buffer pool) to handle frames. The frame size is software configured on a per port basis. These buffers are fixed length buffers large enough to hold the maximum length of the frames and are equally divided for transmit and receive. The receiver does have a higher priority versus other operations associated with the LP card.

With regard to Figure 2-2 when a packet comes into the node from Comm Port 2 the data is placed into an available LP receive buffer via a DMA transfer. It is then moved into the shared memory region of that LP card and now the PPMOS can look at this packet information to extract the required information to determine the destination address and the Logical Channel Number (LCN). The LCN is changed to the destination LCN. LP and port fields are changed to the destination LP and port. The packet is moved to the destination shared memory region. The destination LP polls its shared memory, discovers the packet and copies it into an available transmit buffer and transmits it, via DMA, out on the output communication line (Comm Port 7).

A congestion list queue is also maintained by the PPM. If the destination shared memory is full, at the moment, the packet is placed into this congestion queue until the shared memory can accept more packets. The order of the packets received in the congestion list are preserved. When the destination LP polls its shared memory area to discover that it has a packet, it then copies it into its local transmit buffer and transmits out on the communications line via DMA.

Asynchronous Packet Assembler/Disassembler

The NSN features an integral asynchronous PAD which supports CCITT's X.3, X.28, and X.29 recommendations for connecting asynchronous start-stop devices to a packet switching network. It offers the ability to access any asynchronous host application on the network through a single, low-cost ASCII terminal, workstation, or PC.

The Async PAD supports up to 63 asynchronous ports and one (1) port for the X.25 connection per node for a total of 64-ports; in 4-port increments; with port speeds ranging from 50 to 38,400 bps plus an autobaud capability. When configuring a 64-port node eight (8) LP cards and eight (8) asynchronous Daughter cards are used. The Async only Daughter card physically piggybacks onto the Standard LP card and uses some of the aggregate bandwidth from the LP card. The Daughter card does not plug into a slot in the workstation (ISA bus) but directly onto the LP card. Because the Daughter card has a 62-pin DB connector for connection to the external splitter cable it does use a workstation backplane slot.

Login password protection, mnemonic addressing, dial-up subscriber support, dial-out connections are all available. Also supported are downloadable profiles, custom banner, and logout timer support. Table 2-2 summarizes the subscriber parameters that can be assigned and covers all subscriber types.

On the terminal side, the PAD supports the connection of several devices which are a mixture of terminals, CPU ports, printers, and other Async equipment. On the network side, one or more high-speed ports provide links to the X.25 network. Since the NSNs support a wide range of terminals or other Async devices, individual named profiles can be established for these devices. When you turn your terminal on, you will get a standard screen banner (BANNER.TXT) or your custom banner.

Prompt	Valid Values	Default Values
Subscriber ID	0 - 15 digits	"*"
Required User Data	0 - 12 digits	none
Subscriber Name	0 - 24 characters	none
Node ID	1 - 250	your node
LP ID	0 - 7	0
Port	0 - 7	0
Subscriber Type	1 - 3	1
Group Send Allowed	Y / N	N
Fast Select Accept	Y / N	N
Reverse Charge Accept	Y / N	N
PAD Profile	0 - 10 characters	none
PAD Login Password	0 - 5 characters	none
Preferential CUG	1 - 2 digits	00
CUG	1 - 2 digits	00
Outgoing Access	Y / N	N
Incoming Access	Y / N	N
Hunt Group Node "n"	1 - 250	0
Hunt Group LP "n"	0 - 254	0
Hunt Group Port "n"	0 - 3	0
Hunt Group Algorithm	1 - 3	1
Incoming Calls Barred	Y / N	N
Outgoing Calls Barred	Y / N	N
Network User ID	0 - 2	0
Systematic Redirection	Y / N	N
Redirect on Out of Order	Y / N	N
Redirect on Busy	Y / N	N
"n" Redirected Sub. No.	any valid Subscriber ID	none

Table 2-2 Subscriber Parameters File

Table 2-3 illustrates the X.3 parameters that are available for customizing your profile. The correct setting of these parameters enables the PAD to operate with an attached terminal and your application. When a PAD serves more than one terminal, you can establish and maintain a set of profiles designed for each terminal. The profile name can be up to 10 characters long and the profile description can be up to 30 characters long to uniquely describe each profile. The profile name and description are used in combination when you select the menu option to display existing PAD profile names.

Parameter Reference & Description	Possible Parameter Values
01 PAD Recall Character	0 = Not possible 1 = Possible by character DLE 2 - 126 = Decimal value of ASCII character used to access the X.28 mode
02 Echo	0 = Disable echo 1 = Enable echo
03 Data Forwarding	0 = No data forwarding 1 = Alphanumeric characters 2 = CR 4 = ESC, BEL, ENQ, or ACK 8 = DEL, CAN, or DC2 16 = EOT or ETX 32 = HT, LF, VT, or FF 64 = All other characters in columns 0 & 1 of the ASCII chart. XXX = Any combination
04 Idle Timer Delay	0 = No idle timer 1 - 255 = Inactivity period multiplied by 50 ms
05 Device Flow Control	0 = PAD does use XON/XOFF flow control 1 = PAD uses XON/XOFF flow control in data transfer mode only 2 = PAD uses XON/XOFF flow control in both data transfer & command
06 PAD Service Signal	0 = No PAD messages or service prompt 1 = PAD outputs messages only 4 = PAD outputs service prompt only 5 = PAD outputs prompt and messages
07 Receive Break Signal	0 = No action 2 = Reset packet sent 8 = Escape to X.28 1 = Interrupt packet sent 4 = Indication of break PAD message 16 = Discard output XX = Any combination
08 Discard Output	0 = No 1 = Yes discard output
09 Padding after <CR>	0 = None 1 - 127 = Number of padding characters after CR
10 Line Folding	0 = No line folding 1 - 255 = Number of characters to trigger it
11 Terminal Speed	0 = 110 bps 1 = 134.5 bps 2 = 300 bps 3 = 1200 bps 4 = 600 bps 5 = 75 bps 6 = 150 bps 7 = 1800 bps 8 = 200 bps 9 = 100 bps 10 = 50 bps 11 = 75/1200 bps 12 = 2400 bps 13 = 4800 bps 14 = 9600 bps 15 = 19200 bps 16 = 48000 bps 17 = 56000 bps 18 = 64000 bps 255 = Autobaud
12 Flow Control	0 = Device cannot exercise flow control over PAD 1 = Device can exercise flow control using XON/XOFF characters
13 Line Feed after <CR>	0 = No line feed inserted 1 = LF after CR from remote DTE 2 = LF after CR sent to remote DTE 3 = LF after CR from or to remote DTE 4 = LF after CR when echoing 5 = LF after CR to device when echo remote DTE sends CR 6 = PAD transmits LF when echoing or transmitting CR 7 = PAD transmits LF after all CR
14 Padding after <LF>	0 = None 1 - 127 = Number of nulls characters after LF
15 Editing	0 = No editing 1 = Editing enabled
16 Character Delete	0 - 127 = Decimal value of character used for character delete
17 Line Delete	0 - 127 = Decimal value of character used for line delete
18 Line Display	0 - 127 = Decimal value of character used for line display
19 Edit PAD Service Signals	0 = No editing service signals will be sent 1 = The character \ will be sent for a character deletion. The character xxx CR LF will be sent for a line deletion. 2 = The character BS space BS will be sent for a character deletion. For a line deletion, same sequence performed n times for number of characters. 3 - 126 = The defined ASCII character for a character deletion.
20 Echo Mask	0 = No echo mask 1 = No echo of CR 2 = No echo of LF 4 = No echo of VT, HT, or FF 8 = No echo of BEL or BS 16 = No echo of ESC or ENQ 32 = No echo of ACK, NAK, STX, SOH, EOT, ETB, or ETX 64 = No echo of the editing characters defined by parameters 16, 17, and 18 128 = No echo of characters below the decimal value of 31.
21 Parity Treatment	0 = Disable parity checking & generation 1 = Enable parity checking 3 = Enable parity checking & generation 2 = Enable parity generation
22 Page Wait	0 = Disable page waits 1 - 255 = The number of linefeeds the PAD will transmit before initiating a page wait condition.

Table 2-3 Async PAD Profile File

IBM SNA PAD

The NSNs integrated SNA PAD supports communications between SNA/SDLC terminals, standard cluster controllers, and IBM compatible hosts. The SNA PAD is an interface device that allows SNA hosts and terminals, which do not have X.25 interfaces, to communicate across an X.25 network. The NSNs are compatible with the following IBM protocols and software:

- IBM's SNA/SDLC protocol
- IBM's NCP Packet Switching Interface (NPSI) software
- IBM's Qualified Logical Link Control (QLLC) protocol

Devices attached to SNA PADs may communicate with devices that use integrated X.25 interfaces as long as those interfaces support IBM's standard QLLC protocol. Examples of interfaces that meet this criteria are NPSI on the IBM 3725 or 3745 systems and the X.25 interface on the IBM S/38.

Each port is software configurable to support IBM SNA Host (HPAD) or SNA Terminal PAD (TPAD) side of a sessions, as well as host-to-host sessions (NPAD), and transparent mode (XPAD). XPAD performs transparent PAD operations, that is, it provides the pipeline through which two SNA nodes; using the Q-bit of the X.25 data packet; can transfer SNA control data. The QLLC provides the services of the link layer as perceived by higher levels of SNA architecture. For example, the connection between a PU Type 2 node (SNA 3270/SNA 3770) and a PU Type 4 node with/without supporting the X.25 functions (37x5 with/without NPSI under the NCP). Table 2-4 shows the database and system (default) values that can be established for an SNA HPAD/TPAD configurations.

Prompt	Valid Values	Default Values
Subscriber ID	0 - 15 digits	none
TPAD Number	0 - 1	0
Remote Subscriber ID	any valid ID	none
SDLC Station Address	0 - ff	0
Autocall	Y / N	N
Autocall Timer (minutes)	1 - 255	1
Maximum Retry Autocall	0 - 255	0
PVC	Y / N	N
LCN	1 - 4095	1
XID	0 - 20 hex digits	none

Table 2-4 SNA HPAD/TPAD Configuration

When the SNA port becomes operational, the NSN automatically establishes the virtual connection to the remote destination. The Hunt Group function operates with the PAD ports to provide alternate routing around failures or to distribute calls. The Hunt Group distributes calls from TPADs to available HPADs according to the selected Hunt Group search method. Multiple HPADs are typically established to provide fault tolerance and load distribution.

When an IBM FEP experiences problems, all the remote clusters attached to it are cut off from host applications. Users can select an alternate FEP to establish another connection to the host thereby bypassing the faulty IBM FEP. Since the host destination can be changed dynamically, users can continues to operate. Also, the communications overhead are reduced since the NSNs provide local polling.

The TPAD acts as an entry point for SNA users over an X.25 network. It provides communications between remote SNA Control Units (CU) and IBM host or compatible computers with the NPSI software or connected to the X.25 network through an SNA HPAD. The TPAD supports up to 32 CUs and up to 16 TPADs are supported per HPAD.

The SNA HPAD also acts as an entry point between SNA host computers and an X.25 network. It offers communication between SNA host computers and remote SNA computing devices having X.25/QLLC support or connected to the X.25 network through an SNA TPAD. Each HPAD supports up to 127 CUs over 16 TPADs.

The NSN performs flow control for both the TPAD and HPAD, automatic call setup, and packetizing of SNA information frames. Local polling with the attached devices keeps the network traffic to a minimum.

The SDLC PAD supports connections to SDLC devices in full-duplex or half-duplex, and point-to-point configurations for TPAD connections. Further, you can configure the SDLC frame size up to a maximum frame size of 4,105 bytes as well as the primary and secondary station addresses.

When the SNA PAD is configured for autocall it will automatically place a call when the attached device is powered on. Either the HPAD or the TPAD can be configured to autocall. If HPAD is configured to autocall, when it receives POLL for TPAD, the HPAD will place a call to the TPAD. The TPAD will accept the call if the link has been established with the cluster controller; otherwise, it will clear the call. In case the call was cleared, the HPAD will retry for a specified number of times established by the network administrator. If the call was accepted, then the virtual circuit is established between the HPAD and TPAD, and now the HPAD can start responding to the polls.

NSN Software Architecture

There are three major software modules that comprise AMNET's integrated software architecture corresponding to the three major hardware subsystems:

- Line Processor software
- Packet Processor Module (PPM) software
- Network Management Processor (NMP) software

The first software component is the Line Processor software which controls the operations of the LP card and supplies the interface to the communication ports. It provides all control of the physical port interface (Level I), and executes the logical and link layer (Level II) interfaces. It contains the LAPB (trunk and X.25), the 3X PAD, and the SNA/SDLC PAD applications. Also, there is a real-time operating system developed by AMNET, which is referred to as the Front End Processor Operating System (FEPOS).

FEPOS is the operating system which is loaded during initial program load from the floppy diskette or hard disk. Several applications programs or tasks are also loaded simultaneously, into the LP's on-board memory. When we refer to the FEPOS, we are also taking into account these additional resident tasks on the LP card. FEPOS controls these tasks on a real-time basis, provides task dispatching, memory management, and I/O and interrupt handling.

The second software component is called the Packet Processor Module Operating System (PPMOS) which is resident in the workstation. The PPMOS carries out the routing, addressing, flow and congestion control, priority determination, packet switching, and provides the direct user interface capabilities. PPMOS uses an efficient round robin algorithm to dispatch tasks. All common data structures are accessed through operating system calls for efficiency. And like the FEPOS, it also uses the shared memory scheme.

The Network Management Processor (NMP) software consists of an operating system, various utilities, and application programs. The operating system and its associated utilities provide a controlled environment in which the network management application programs operate. The N7900 Network Manager programs allow the operator to update the database, examine statistical information, change network parameters, start up and shut down the system, and selectively alter functions at the LP level. All the N7900 operations are covered in greater detail in Section 3 of this manual.

Front End Processor Operating System

The Front End Processor Operating System (*FEPOS*) and a set of tasks are resident in each LP card. *FEPOS* is real-time executive to coordinate the many concurrent activities (tasks) within the LP application system. The operating system is a pre-emptive operating system based on a fixed priority scheme for each task. *FEPOS* implements functions necessary to maintain the performance oriented system required by the LP applications.

FEPOS maintains a set of attributes for each task in the system including the priority and execution state of the task. Priorities range from the highest to the lowest and it performs the following functions:

- Initialization
- Scheduling and dispatching tasks
- Buffer allocation and management
- Timers
- Inter-task communications and synchronization
- Message management

The *FEPOS* real-time system has the ability to respond to events occurring asynchronously. Timer interrupts provide the mechanism to support this real-time environment. A timer is activated every 20 milliseconds. After the servicing of the interrupt, a task can request that a message be sent to another task or to itself. This provides for inter-task communications and task synchronization.

This inter-task communications is handled via a queue. Messages (both normal and priority messages) are typically placed in the queue and are retrieved on a first-in-first-out basis. Some messages can set a priority flag and these messages will be placed at the head of the queue. Normal messages will be added to the tail of the queue.

Tasks can be in any one of four (4) states: Ready, Wait, Running, and Suspended. This is the basis of any real-time operating system. When a task is executing, it is in the *Run* state. Any task that is eligible to run is in the *Ready* state. A task is in the *Wait* state when it expects a message at its queue and the message has not arrived. This task is now a candidate to be *Suspended.* There are several ways a task can be *Suspended.* It can suspend itself, call a higher priority task to execute, or call a task's "timed wait" times out. Additionally, tasks can request buffers dynamically and the *FEPOS* also manages this system resource.

LP tasks controlled by the *FEPOS* are illustrated in Figure 2-3. Following the figure, each task and the function it performs is briefly explained. The tasks are arranged according to their actual priority structure in the LP environment with the highest priority task first.

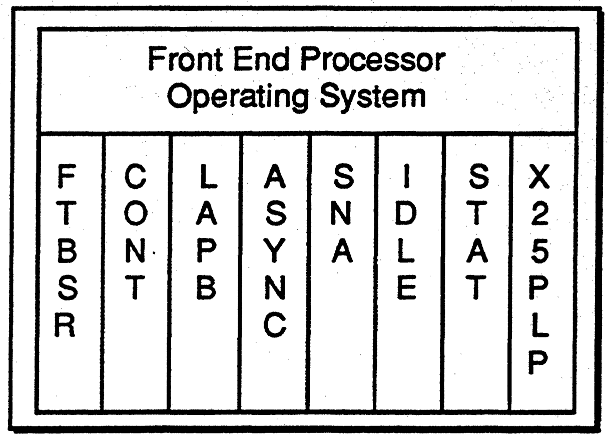

Figure 2-3 LP Tasks

FTBSR — This task acts as the interface between the LP and the PPM (i.e., waits for the messages to be sent to the Node Buffer Memory). It performs the polling of the shared memory region, flow control, and the management of its buffers. The flow control mechanism will stop writing to the LP shared memory region. *FTBSR* is the highest priority task in the LP system and is called when the Watchdog Timer expires.

CONT — This is the LP Controller task. Its function is to process messages sent to it by the Node Buffer Memory.

STAT — This is the third high priority task that collects statistics from the other tasks.

LAPB — It performs the X.25 Level II LAPB functions.

X.25 PLP — This task provides the X.25 DTE packet level procedures to the PADs including SNA/SDLC, Async, X.25, X.29, etc.

SNA — All SNA/SDLC PAD functions are handled by this task.

ASYNC — This task accomplishes the 3X PAD (X.3, X.28, X.29) functions.

IDLE — When no other task needs control, the *IDLE* task is then executed. It is the lowest priority task but performs two (2) needed functions: first, time computation; and second, sending a wake up message to *FTBSR* if there is a message waiting in the shared memory region.

FEPOS Initialization

Before the LP application tasks can begin operation, each task, its data area, and the LP resources controlling it must be initialized. Upon system startup or restart, after system diagnostics or initial program load procedures have run, the LP initialization routine (*FEPINIT*) establishes a "known" execution environment for all the tasks.

The initialization routine first zeros out each task's data and program area and then establishes the tasks stack area. Next, the remaining free memory in the LP is divided into individual buffers. These buffers are then linked to form the Free Buffer Queue. The Free Buffer Queue is used by both the LP and PPM. Finally *FEPINIT* places all tasks in the Ready state.

PPM Operating System

PPMOS is a multitasking operating system which supports task scheduling (dispatching) and memory management operations. The operating system is resident in the workstation or the node. All the tasks (processes) that run in this environment are established using a priority scheme. Higher priority tasks have a longer execution time than tasks with a lower priority. The scheduler implements a non-preemptive round robin scheme on a per task basis. This means that the task, once started, has control of the node until it relinquishes control. The PPM operating system cycles through this list of tasks in a pre-defined order and based on a set priority.

PPMOS and all of its application programs are automatically downloaded into the nodes local memory from the local boot disk during normal startup operation or after a restart sequence initiated by the operator. Figure 2-4 illustrates the tasks controlled by the PPM Operating System. In this section we will cover a majority of the tasks in the PPM with the noted exception of the Computer Interface Process (CIP). This is covered in greater detail in Section 5.

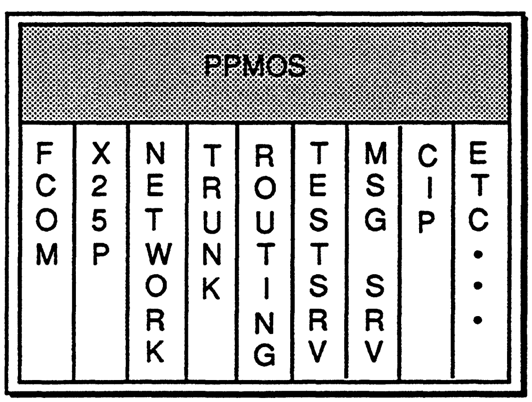

Figure 2-4 PPM Tasks

PPMOS supports both intra-task and inter-task communications. With intra-task communications tasks can send messages to themselves. Inter-task communications provide a communications scheme in which different tasks can exchange messages. Both communications schemes are supported via queues. The messages sent and received can be normal or have a priority associated with the message. Priority messages are placed at the top of the queue whereas normal messages are placed at the bottom of the queue. All the queues operate in a First-In-First-Out (FIFO) basis.

The memory management supports dynamic allocation of variable-length buffers from a Common Memory pool. The overall size of the Common Memory region is approximately 192 Kbytes. The Common Memory module supports other functions in addition to the intra- and inter-task communications. It holds transient connection record information, flags and pointers to provide a global system data repository available to any task on the workstation. All tasks are establish (linked) with this Common Memory region, thereby, providing this global systems area.

At the beginning of the Common Memory region the following system information parameters are established:

- System Process Directory - During initialization each task establishes its queue for passing messages. This directory is a pointer to a table that is indexed by task ID where this task can normally receive data.
- System Network Record - This is a pointer to an in-core structure that contains network information.
- System Node Record - This is a pointer to an in-core structure that contains node information.
- System Line Directory - This is a pointer to a table that is indexed by the port number. Each entry contains a pointer to a port record that was downloaded from the Network Manager.
- System FEP Status Table - This table entry provides a pointer to a table indexed by LP number to indicate whether or not the LP is active.
- Test Services Line Table - This table entry is a pointer to a line directory entry for each LCN created by Test Services.

All task requests are processed via system calls. A system call provides an elegant means of satisfying the task requirements for requesting information, passing arguments, and retrieving data. One of the most frequently used system calls is to read the system parameters. This will return the total number of queues in the system and number in use, the total number of 16-byte buffers available and in use, number of pool buffers used during a one hour window, as well as the total time spent in PPMOS.

Some of the processes we will look at are the FEP Control Process, X.25 Level III Process, Routing Calculation Process, Trunk Process, and the Network Process. Some of the other services as CIP and the test services are covered in greater detail in the other sections of this document.

Network Process Manager

The Network Process Manager (NPM) process performs the following functions:

- Receives and distributes messages from the N7900 Network Manager
- Processes messages from the Network Manager
- Sends messages to any locally attached Network Manager
- Sends messages to another NPM process on a node that has a Network Manager
- Performs the re-IPL operation on the node
- Reboot the Network Manager
- Monitors the locally attached Network Managers status

The NPM acts as the link between the PPM process and the network. The NPM process provides the capabilities for a PPM process in one node to send a message to another process on the same node or to a remote node. Messages are processed from various sources and come from various sources. For example, the N7900 Network Manager places its message in the Common Memory buffer as well as the other processes controlled by the PPMOS. Messages from the LPs or from the network would normally be placed into the Node Buffer Memory region and the NPM process handles all of them.

The NPM process receives status indication messages from the network to maintain a node status table. This node status table shows the status of each possible node and may even contain a list of messages destined for that node. The possible node status indications are: connection down, connection up, or flow control off. When the last message is processed data is resumed being sent to the node once it detects the RR.

Trunk Process

The Trunk Process provides the Level III Packet processing for both dedicated (leased) and dial-out (switched) trunk lines. For dial-out trunks, it also controls dial-out and hang-up processing. Under normal conditions a node and its trunk lines are operational within 65 seconds after powered-on.

The normal transitions for a dedicated trunk from a power-up condition is as follows: first, Trunk Fail (from power-up or re-IPL state); second, Link Disconnect (this occurs when DSR, CTS, RTS, and Carrier Detected are seen); third, Link Up (SABM, UA sequence complete); and fourth, Link Established (node IDs have been exchanged). The

Trunk Process controls the transition from Link Up to Link Established. Once the trunk is established, it is available for data transmission.

FEP Control Process

The FEP Control Process has two main functions. It performs the operator initiated functions and handles the LP bound messages from the Network Manager via the NPM and CIP processes. The main job of the FEP process is to handle operator initiated functions which require LP action as:

- Enable an LP port
- Disable an LP port
- Examine port or LP parameters
- Change port or LP parameters

The FEP Control Process is also responsible for creation and maintenance of the network, node, and port records in the nodes common memory region. Also, it transmits, modifies, and examines the control messages to the LP addressed, and retransmits them until an ACK is received, or until the retry count is reached.

During the node start up process, the FEP Control Process receives the node record message (number of trunk and data ports, maximum line number) from an Network Manager which has been routed through the local copy of NPM. Then the LP receives the network record which includes such information as node number, packet and window size, default network window size, reset timeout limit, and retransmit limits. Next the port record message is received, one for each active port. The formats are different for each type but would include the port number, throughput class, port type, LP type, and physical interface.

All the parameter messages are now received from the N7900. Each message is relayed to the destination LP with an acknowledgement expected within eight (8) seconds. When the LPs acknowledge all of the aforementioned messages they are marked as active in the system.

X.25 Level III Process

The X.25 Level III Process provides and maintains the interface for PVC and SVC services between the DCE and the DTE. Implementation of the X.25 Level III process is based on the packet level operations of the CCITT X.25 Recommendations. This process is supported by a myriad of sub-tasks that handles the following:

- Packet Level Flow Control
- Packet types and formats

- Diagnostic and message codes
- Call Setup, Call Clear, and Call Verification
- Group Send ACK
- Create and remove PVC requests
- Packet sequencing

To facilitate node access the X.25 process maintains two (2) separate directories, namely; a connection directory and a line directory. The connection directory is a link list of the information on all channels, keyed by their associated connection number. Almost all packet messages that are exchanged between the Level III process and the network process use the connection directory. The line directory is a link list of all the information on all channels on each port and is maintained by the system via its port record file.

Supporting messages is one of the major functions of the X.25 Level III process. Internal line status message codes generated by DCE from the remote end would include restart port request, channel enable request, create and remove PVC/SVC requests. Also internal packet messages are generated as data, RR, RNR, and reset request. Some external packet messages generated by the DTE from the FEP are: call and clear request, reject, reset request and confirmation, data, receiver ready, and receiver not ready.

The X.25 Level III process also provides us with diagnostic codes. These codes are used by the network operator to analyze error situations that could arise. When we established the diagnostic codes, we considered them from the viewpoint of cause and effect. Therefore, the diagnostic codes are broken down into four categories: clearing causes, reset cause, restart cause, and network generated diagnostic codes.

In order to understand the cause/effect relationship lets look at some examples. For a clearing cause we could get number busy, out of order, access barred, and RPOA out of order. A reset cause would be generated by remote procedure error, local procedure error, or remote DTE operational. Restart causes are local procedure error, network congestion, and network operational. Most of the network generated diagnostic codes deal with the invalid packet type, virtual circuits, and facility codes (i.e., invalid p(s), invalid p(r), timer expired, facility code not allowed).

NSN Windows

There are three (3) kinds of windows supported in the N7400 and N7500 environments:

- Packet Level Window
- Frame Level Window
- Network Level Window

You have control of all three (3) level windows for your network configuration. Figure 2-5 represents three (3) switching node on the network. Each connection between nodes has an associated Frame Window. The Frame Level Window size can be set from 1 to 127 with a default value of 7. On the N7400 and N7500, the Packet Window is local to the link that is involved in the subscriber connection. The last window supported on the backbone network is the Network Window. The Network Window handles the end-to-end control from the source node to the destination node.

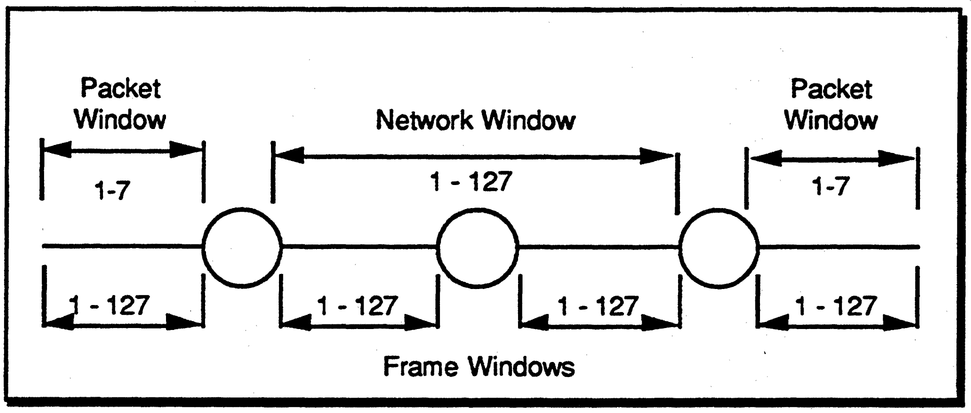

Figure 2-5 Network Windows

Depending on the network configuration there could be a large delay between the source and destination nodes. With a relatively large Network Window this would allow the subscriber to keep interjecting packets into the source node and into the network until the Network Window is full. This provides for an efficient and effective window control mechanism for your end-to-end transmissions.

Also on the NSNs, we have incorporated a NBM region or packet buffer area. Figure 2-6 illustrates how this is implemented on the NSNs. When a packet enters Comm Port 9 it is copied from the LP receiver buffer region to the shared memory area of LP #2. If this is an X.25 frame or subscriber, we will hold a copy of it in the LP receiver buffer region until we get an acknowledgement. We also hold it in the NBM memory just in case we have to retransmit it at a later time due to line noise.

With a copy of the packet in the shared memory region, we now can break it up and take the header portion and move it into the NBM memory. The header portion is all that the PPM requires. The data portion remains in shared memory. After the PPM determines what is the best route, the header information is copied from the NBM, and the data portion from the shared memory over to the transmit shared memory region. In our example this happens to be LP #5. Once it is in the shared memory region of LP #5 it is then copied to the transmit buffer area and then ultimately out on Comm Port 21.

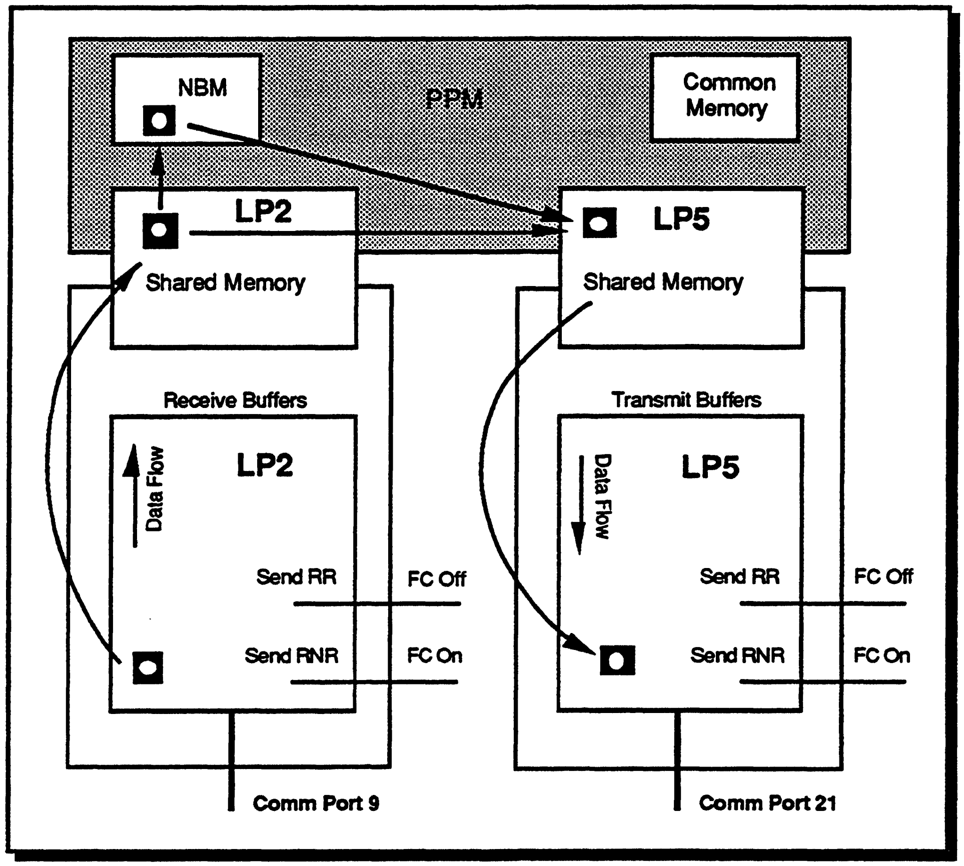

Figure 2-6 Node Buffer Memory

Once this packet gets transmitted out over the network the intermediate buffers are released (i.e., shared memory and transmit buffer), and when a packet level acknowledgement comes back we return (free up) the receive buffer and NBM. The PPM maintains a list of all the buffers that can be released and returned to the free buffer pool. The PPM sends a message down to the LP card via the shared memory interface and returns a group of these buffers at once as well as releasing the ones previously allocated in the NBM.

NSNs Packet Level Flow Control

As we have previously stated, you have control of all of these three level windows (see Figure 2-5) in your system. Therefore, you can make adjustments to some of them, none of them, or all of them. But what happens when the packets do not reach their destination? Without the acknowledgement do we release the NBM buffers or receiver buffers? Will the Node Buffer Memory fill up? Will the congestion queue fill? Or both?

We know that we can adjust the Network Window depending on the number of packet buffers in use and this actually comes from the NBM - they are synonymous. This adjustment is balanced by a related amount of common memory in use. Common memory is where all connections records are kept and is also an integral part of the PPM. Therefore, if we have a lot of connections or common memory in use we will tend to make our adjustment quicker than if we have just a few connections. The reasons is that with a lot of connections in use we could get a burst from the subscribers which could flood memory.

As packet buffers in use get above 30%, we begin to adjust the size of the Network Window. It starts out at some maximum size either specified by the operator or the system default value which is 20. We start adjusting the window downwards. As long as the packet buffers in use stay above 30%, we continue to adjust it downwardly. In fact, we adjust it downwards rather drastically. For example, we cut it in half each time we make the adjustment. As long as the packet buffers in use stay between 30% and 60% we will not take it down below 2. Two (2) is the smallest increment that we will adjust to unless the packet buffers in use goes above 60%. If this occurs then we will adjust the Network Window down to 1.

The Network Window is a mechanism used to transmit packets without an acknowledgement. We will still receive any number of packets and acknowledge them as they come in. To make things balance the window size that the node is using is included in the routing update packets that are transmitted around the network. Therefore, we use the smaller of the local window or the window that the node is using when we transmit to the node.

Remember, the Network Window controls a call on a per connection basis and if one connection has its window full that doesn't prevent another connection from sending until its Network Windows reaches its maximum value. This provides additional flexibility by making a per connection adjustment rather than making it for all connections on the given node, reducing the number of packets that have to be held awaiting retransmission, and protecting the buffer utilization (packet buffers in use).

We mentioned that there are other windows that can be adjusted, namely; the Packet and Frame Window. Can we do anything with them? X.25 doesn't allow us to adjust the Packet Window size. We know that the window size is set to seven (7). That means that the subscriber can send us seven (7) packets, and we can acknowledge all seven (7) packet at once, and then the subscriber can send us another seven (7) packets. But what happens if we only acknowledged a certain number of the packets?

When the buffer utilization starts to increase we can acknowledge a fewer number of the packets at that point in time. For example, if there are seven (7) unacknowledged packets we could acknowledge two (2) of them instead of all seven (7). That means that the subscriber can now only send us two (2) more packets. That allows us to effectively control the Packet Window size. This means that we will have some packets just sitting there but we can't get a burst of traffic since we control the acknowledgements. Over a period of time this will reduce the buffer utilization or packet buffers in use.

And the last resource is the Frame Window adjustment via flow control mechanism (i.e., RNR). The Congestion Queue is an essential part of the PPM (see Figure 2-7) and is utilized as a last ditch effort. Each LP card has a fixed number of buffers dedicated to support the receiver and transmit areas. The receiver and transmit buffer allocation is about equal with some additional buffers reserved for the receiver side. In most cases, these buffers are fixed length and they are big enough to hold the maximum length of the frame.

When the utilization of the transmit buffer for Comm Port 21 reaches its threshold it sets a flag for its shared memory region not to accept any more packets. Once this occurs, packets are now placed into the Congestion Queue. If this queue gets too large then LP # 2 shared memory region will start to fill followed by an increase in the receiver buffer utilization. When this occurs a RNR message will be sent informing the X.25 process to stop acknowledging packets. This is the ability of the destination to quench the source.

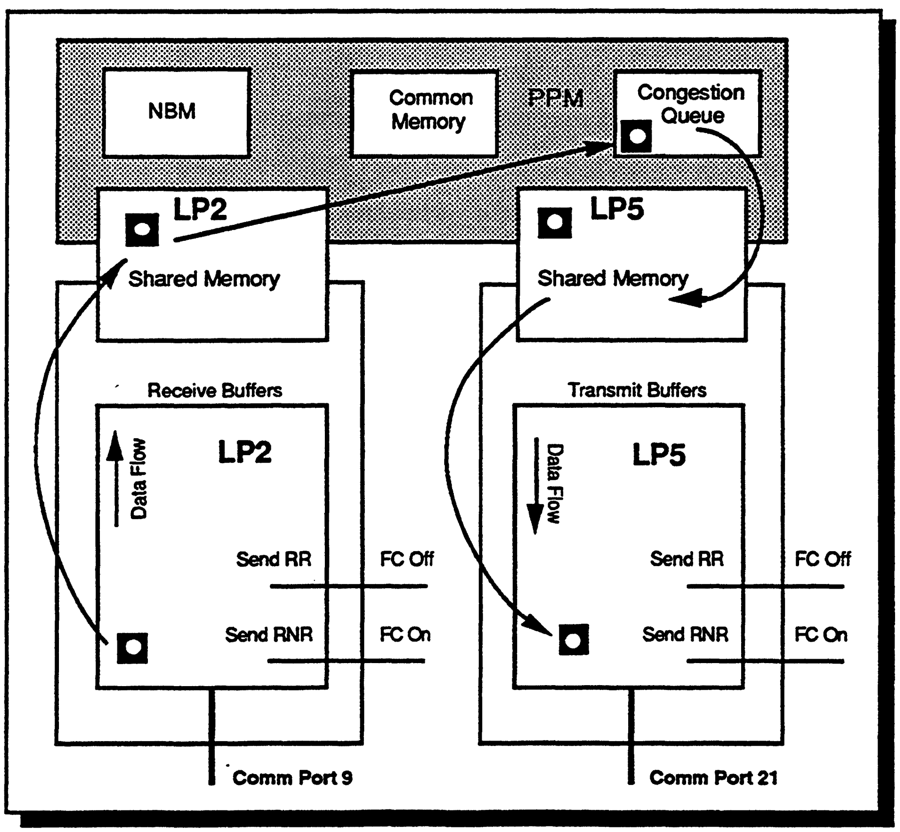

Figure 2-7 Congestion Queue

Conversely, as the buffer utilization drops below 10%, we automatically begin to adjust the Network Window back to its original value. But we will adjust the window back to its original value slowly, that is, one (1) increment every 20 to 25 seconds. It goes down quickly but is restored slowly which is the typical scenario in a congestion situation.

AMNET's Routing Calculation Process

AMNET uses a proprietary adaptive routing algorithm to perform packet switching operations. Routing of packets is governed by the routing tables which are stored in each node. These routing tables are not specified in the network configuration. Instead, they are built by each node as the node is brought on-line and the node is able to discover who are its neighbors. Thereafter, each node updates its routing table to reflect congestion or changes detected in the network topology. Changes to the routing table immediately alters the routing traffic through the node.

The routing algorithm minimizes delay while maximizing throughput. It is tuned to select the shortest path (in terms of delay) for every packet traversing the network. As a packet travels through the network the routing tables at each intermediate node are consulted to determine the best trunk over which to forward the packet. The network utilizes any suitably uncongested path in delivering data.

The routing algorithm provides reliable network operation by automatically routing traffic around failed nodes or trunks as soon as a failure is detected. Additionally, AMNET's End-to-End protocol ensures the delivery of packets in order by re-sequencing any out-of-order packets at the destination node. This mechanism prevents data loss by buffering packets at the source node and retransmitting any unacknowledged packets.

Shortest Path First Algorithm

AMNET's Shortest Path First (SPF) algorithm is a revision from the original Arpanet algorithm. The older Arpanet algorithm has significant problems in regards to response time and looping problems. The three biggest problems encountered using this algorithm are: first, it is slow to respond to significant network changes but is too quick to respond to trivial changes; second, looping between adjacent nodes can last for several seconds; and third, the routing packet update size is related to the size of the network (i.e., the bigger the network the bigger the routing packet update size).

By using the Shortest Path First algorithm almost all loops are eliminated or the loop is very short lived. Another advantage is that it is very responsive to changes in the network condition.

For example, the NSNs will change immediately on a link failure. Also incorporated in our algorithm is that the routing packet update size is not related to the size of the network. The SPF algorithm generates a lot of packets but uses a small packet size, therefore, subscriber data packets are not blocked from being transmitted.

The Shortest Path First algorithm has three major components:

- Delay Measurement
- Propagation of Delay Information (Routing Update)
- Routing Calculation

Delay Measurement

The standard of measurement used in the N7400 and N7500 Network Switches is delay measurement. There are three (3) parts to delay measurement through a node in our network:

- Queueing Delay
- Transmission Delay
- Propagation Delay

Queueing Delay (see Figure 2-8) is composed of queueing and switching delays associated with the local node. It is the total amount of time that the packet remains inside the packet switch. We measure the Queueing Delay at the node by placing a time stamp when the last bit in of the packet enters the node. This packet is then processed and switched in the PPM and we initiate another time stamp as the first bit goes out of the node. For the actual Queueing Delay we then subtract the time stamps of the first bit out minus the last bit in. This provides us with the actual amount of time that the packet has spent inside the node.

The PPM maintains a clock in the shared memory area. This clock increments every 20 milliseconds and all the LPs on the same node work off this synchronized clock. Therefore, the time stamp measurements at Comm Port 9 relates directly to the time stamp at Comm Port 21.

The second component is Transmission Delay which is the amount of time it takes to get the frame out onto the line. This is the bit rate multiplied by the number of bits in the frame.

And the last component is the Propagation Delay. That is, the time it takes for the bits to move from one end of the line to the other. On a terrestrial line this is negligible but on a satellite line this could be very large. Currently, on the N7400 and N7500 we ignore the Propagation Delay. The Propagation Delay is from one node to another node and since the nodes are not time synchronized it is difficult to measure.

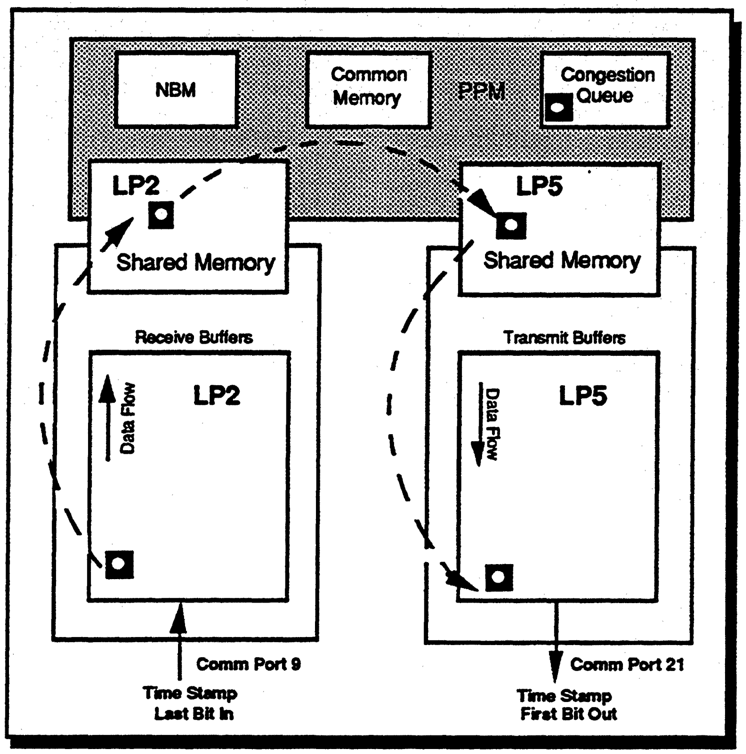

Figure 2-8 Queueing Delay

Routing Update Process

In a stable network each node generates routing updates once every 60 seconds. Conversely, in an unstable environment this could occur every ten (10) seconds or immediately on a trunk failure. This *Routing Update Process* is performed because each node in the network maintains its own database describing the entire network topology and the line delays. Using this database, each node independently calculates the best path to every other node, and outgoing packets are routed accordingly.

There are several steps required for this to be accomplished. First, we have to monitor both the line status and line delay data. Changes in the line status, such as trunk down, and trunk established are communicated immediately to the *Routing Update Process*. Queueing delay information is collected (time stamped) on a packet-by-packet basis every 20 milliseconds. This queueing delay data is then averaged and monitored every ten (10) seconds and is used as a running total of the change in delay on a line-by-line basis.

There is a separate average (one average per trunk) for every neighbor of a node. For example, if we look at Figure 2-9 we can see that the Hartford node has three (3) neighbors (i.e., Boston, NYC, and London). Therefore, it would have three (3) averages for each trunk. If there were multiple trunks to a single neighbor then we would only save the fastest one to that neighbor.

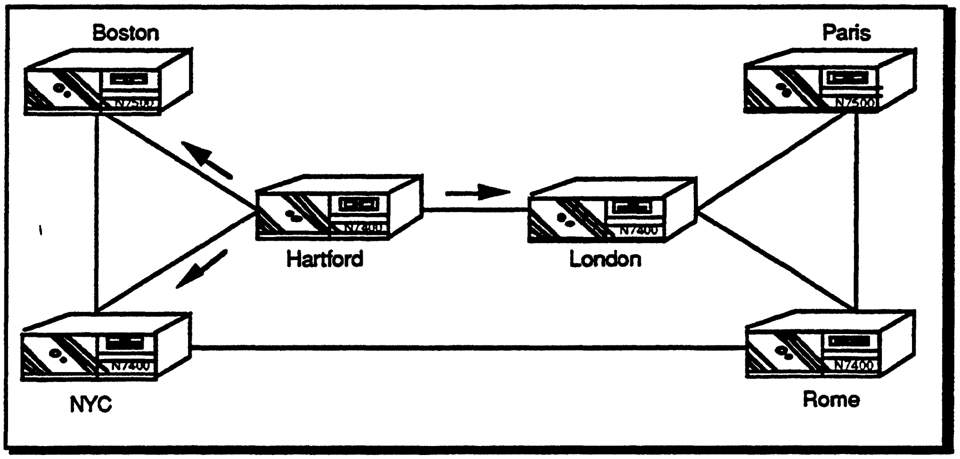

Figure 2-9 Routing Update

Now, to the Queueing Delay averages we add the weighted transmission delay for a 256 byte packet. We select a known packet size because the measurements have to relate to each other. For a typical node, the transmission delay is by far the larger component of delay. Basically, the 256 byte size packet reflects the fact that a majority of times we could end up with a zero queueing delay. For example, if you have zero queueing delay looking at a small packet through a 1,200 bps line it may not appear any different than a 4,096 byte packet over a 56 Kbps trunk line.

To get the weighted factor we use a 256 byte packet and we multiple it by the trunk line speed. Therefore, if one line is operating at 56 Kbps we would take 256 and multiple it by 56 Kbps to yield 14.3 million. Conversely, if it was a 9,600 bps line the weighted factor would be 2.5 million. Now the differences between a fast and slow line shows up dramatically; 14.3 million for a 56 Kbps line versus 2.5 million, respectively.

With the current delay now available the *Routing Update Process* determines whether the change exceeds an allowable threshold. If it does it then disseminates a routing update packet containing the new information. This allowable threshold is not a constant; instead it is a decreasing function of time. Whenever there is a large change in the current delay, it is desirable to report the new delay as soon as possible, so that routing can adapt quickly. But when the delay changes are only slight it is not important to report it quickly, since it is not likely to result in important routing changes.

There are two time factors that come into play for this routing update packet. First, every 60 seconds we will unconditionally generate a routing update packet; and second, is the allowable threshold which is initially established at milliseconds. If the current delay has exceeded a certain threshold since the last routing update packet we begin a new routing update cycle. If the change is not greater than the threshold, then twelve (12) milliseconds is subtracted from the threshold. So the next threshold is 48 milliseconds. After another ten (10) seconds has elapsed we perform this operation again because the *Routing Update Process* monitors the current delay value every ten (10) seconds. And if it is not greater than 48 milliseconds, then it is decreased to 36 milliseconds. And so on, and so on, and so on; until either the 60 seconds has elapsed or the threshold equals zero. During normal operations both conditions will occur at the same time.

The *Routing Update Process* sends the routing update packet to the NBM for distribution. The routing update packet contains a list of immediate neighbors and delays to these neighbors. Routing update packets are distributed by flooding them through the network. That is, each node sends each new routing update packet it receives on all its lines. The same packet is sent on all lines, and no addressing information is required. Flooding is independent of the routing algorithm, this makes it a safe and reliable scheme. Routing update packets are transmitted over all lines, including the input trunk line. The echo over the input line serves as an explicit acknowledgement to the sender.

For example, the Hartford node (refer to Figure 2-9) takes the routing update packet and sends it out on its three (3) trunks. The Boston node gets the routing update packet and sends it out on all of its trunks including the one it came in on. This basically serves as an acknowledgement to the Hartford node as well as forwarding it onto NYC. The NYC node sends it back to Boston, to Hartford, and onto Rome. Every routing update packet must be acknowledged.

In summary, each node rebuilds its routing tables (routing directory) from all routing packets received. These routing update packets happen asynchronously and randomly or every 60 seconds. Every routing update packet has to be acknowledged.

Routing Calculation

The final stage is the routing calculation and it is called whenever there is a change in the network topology or whenever a significant change occurs in the current delay. The routing calculations use the network topology and the line delay data to calculate the *'best path'* to every other node. The routing calculation uses the SPF algorithm which generates a topology tree that represents the minimum delay paths from a specific node to every other network node. Then a routing table is built by traversing the tree of minimum delay.

We will use Figure 2-10 to illustrate how this tree is built and then how the final routing table is generated from the tree. The units of measurement on Figure 2-10 represent delay time. What is illustrated is the current delay from node one to another node. Each unit should be multiplied by 2.5 milliseconds for an actual representation. Therefore, the delay from the NYC node to Boston is ten (10) units or 25 milliseconds. Note that the delay in one direction is not related at all to the delay in the other direction.

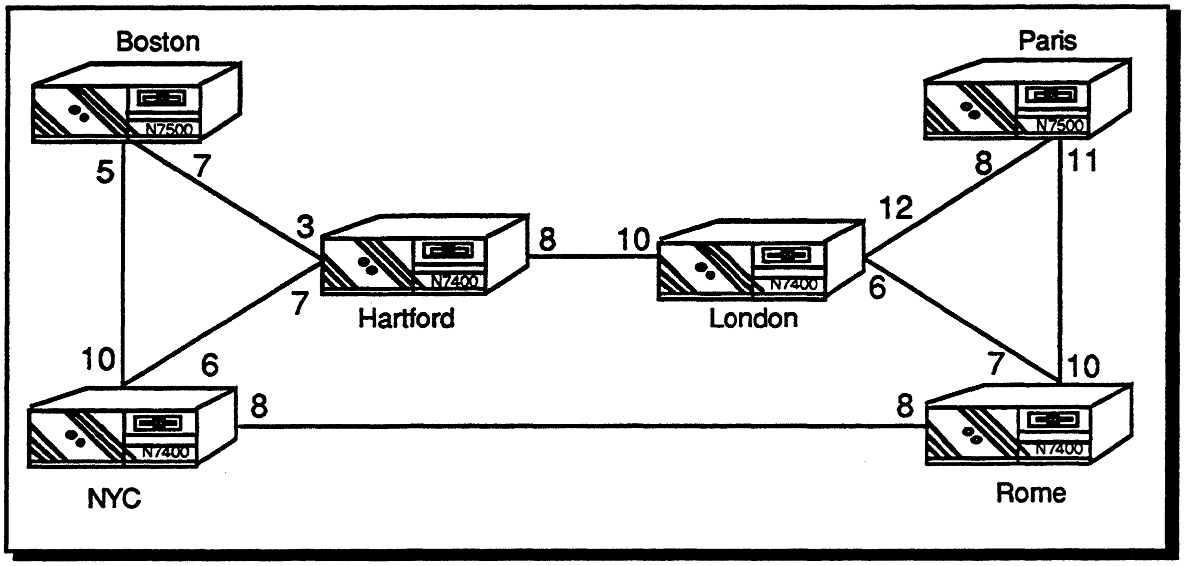

Figure 2-10 Routing Calculation

The routing update packet provides us with the following information:

- All the nodes in the network
- All the links between the nodes
- Current delays between nodes

The reason it is called the Shortest Path First algorithm is from the way that the tree is constructed. Therefore, we will build the tree from Boston out into the network. If you refer to Figure 2-10 the quickest node to get to from Boston is NYC. Place the Boston node at the root of the tree and add the NYC node under the tree with a delay of five (5) units. The next quickest node is Hartford which has a delay of seven (7) units or 17.5 milliseconds. Add Hartford to the tree with a delay of seven (7).

Remember that Boston is the root and we are trying to get to the next quickest node on the network. The Rome node is accessible either from going through the Hartford and LA nodes or directly from the NYC node. If we look at the units of delay from Boston to NYC we see that it is five (5) units. The delay from NYC to Rome is eight (8) units for a total of 13 units.

If we look at Boston going through the Hartford and London nodes the delay would be seven (7) units plus eight (8) units plus six (6) units for a total of 21 units. Therefore, we add Rome to the tree, under NYC since the delay is only eight (8) units. Now the next quickest node is London and it comes off of the Hartford node. And finally Paris. If you perform the calculations the quickest way to get to Paris is via Boston, NYC, and Rome with a total unit delay of 23 units versus 27 units going through Boston, Hartford, and the LA nodes. Now all the nodes in our network are listed in our tree (refer to Figure 2-11).

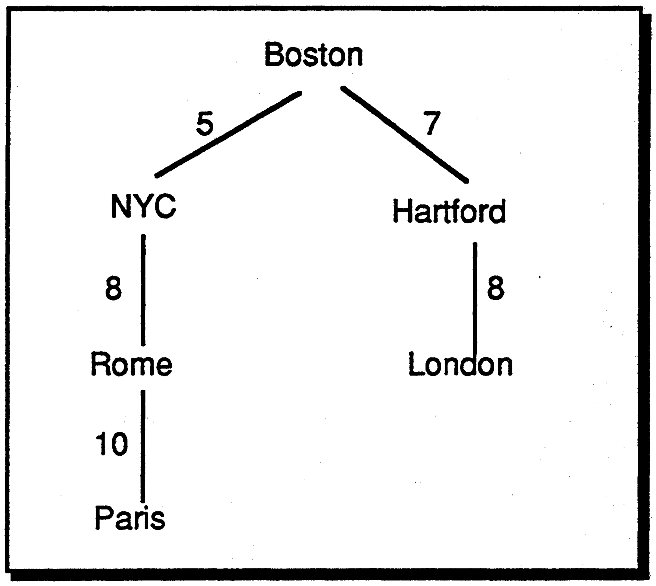

Figure 2-11 Network Topology Tree

Now we can construct our routing table from the topology tree. The routing table is used by each node to indicate the neighboring nodes that could be used to route a packet to its final destination. For our network we want to determine our neighbor, the best way to get to this neighbor, and the number of hops required. This is accomplished by working our way down the topology tree and constructing our routing table. This is illustrated in Table 2-5.

The first column in the table is the node ID, followed by the best path to get to our destination, and finally the number of hops required. Since Boston is our originating node nothing is required. The next hop down is NYC and we get there directly by going to the NYC node. This also holds true for Hartford. Both of these are only one (1) hop away. To get to Rome requires two (2) hops going through NYC. Now for London we have to go through the Hartford node. The number of hops to accomplish this is two (2). Paris requires three (3) hops. We would have to go through NYC, Rome, and finally Paris. The Boston node can now refer to this routing table to make its routing decisions.

Node ID	Best Path	# of Hops
Boston	-	-
NYC	NYC	1
Hartford	Hartford	1
Rome	NYC	2
London	Hartford	2
Paris	NYC	3

Table 2-5 Network Routing Table

Routing Update Load

The routing update packets do impose some load on the network. Each trunk has one packet in and one packet out, in a stable network each node generates a routing update once every 60 seconds.

Now the load on a given node in packets per second is calculated by the formula

N*T / 60

where N = Number of nodes in the network
T = Number of trunks on this node

For example, we want to calculate the load at the Hartford node (refer to Figure 2-10). The total number of nodes in our network is six (6). The number of trunks on the Hartford node is three (3). Therefore, we would get 6*3 or 18 over 60 (18/60 or 3/10) for 0.3 packets per second. Lets expand this and look at a network that has 100 nodes with the same three (3) trunks we would get the following: 300/60 or 5 packets per second. Therefore, as the number of nodes increases so does the number of update routing packets throughout the network.

SECTION 3

NUCLEUS 7900

NETWORK MANAGER

The Nucleus 7900 Network Manager

Effectively managing today's complex data communications networks requires a comprehensive, yet user-friendly network management system like AMNET's Nucleus 7900 (N7900) Network Manager. The N7900 Network Manager meets the needs of today's network administrators and operators by providing powerful functionality and features like:

- Centralized or distributed management support
- Icon based, color graphics interface
- Comprehensive statistics and billing
- Real-time alarms and event reporting
- Customized alarms
- Menu-driven configuration management
- Multiple levels of password security

The N7900 incorporates a graphical user interface which is consistent across all functions, including monitoring, configurating, and administrating of the network and its resources. Immediate visual indication of the location and severity of alarms allow the operator to respond rapidly and efficiently to changing network conditions. With the visual presentation available on the N7900 the operator becomes familiar with a geographically dispersed network.

The N7900 is menu-driven allowing the operator to perform on-line control, configuration management, node monitoring, statistics operations, and performance measurements from the same workstation. Operator training requirements are reduced through the use of the menu-driven interface and graphical icon screens.

The N7900 supports the network configuration database, call setup, call accounting information, and continuously collects data on the status of every device and transmission circuit in the network. Status information including utilization, error rates, and current configuration. Real-time EIA data set signals are all available. Also, the operator can solicit the node to find out what is its current configuration, ports can be taken in and out of service, diagnostics performed, thresholds changed, and test services performed using the capabilities of our N7900 Network Manager.

Modifications such as changes to trunk speed are performed on-line with no impact to other ports in the network. These on-line changes are temporary and are not reflected in the master database. Therefore, if the node was re-IPLed it would revert back to its prior configuration.

N7900 Hardware Structure

The N7900 Network Manager is based on an industry standard 386 based workstation, color display monitor, 8 Mbytes of RAM, an 80 MB hard disk, floppy disk drive, serial printer, and a Parallel Interface Adapter (PIA). The N7900 uses the powerful UNIX operating system. The serial printer is the default error logging device on the system where all alarms and events are logged as well as being automatically archived on the system disk.

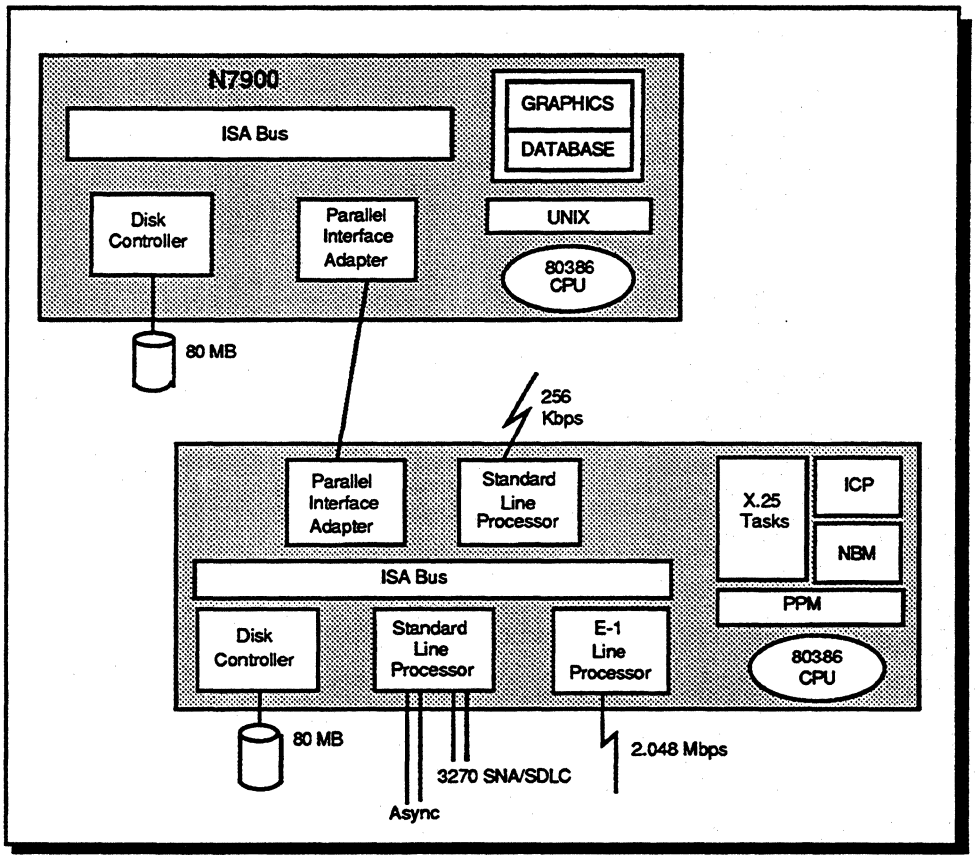

Figure 3-1 N7900 Connection

A PIA card and cable are connected to an adjacent N7400 or N7500 so that the N7900 can access the network. Figure 3-1 illustrates a N7900 connected to an N7400 Network Switch. The high-speed PIA interface provides the network connection into the X.25 domain. The color display monitor is used to configure the network, build and update the network database, change on-line port parameters, customize alarms and internodal trunk capabilities and provide the direct operators interface.

N7900 Software Architecture

The N7900 software consists of the operating system, various utilities, and a myriad of application programs or tasks. The operating system and its associated utilities provide the control environment in which the network management application programs operate. During normal operations the UNIX operating system is transparent to the network operator.

Figure 3-2 illustrates some of the applications and tasks in the N7900. In the following sections we will briefly highlight a portion of these tasks to provide you with a better understanding of their capabilities and some of the capabilities of the N7900 Network Manager. (Programmers, system designers and system integrators that require greater detail can reference the *LP and PPM Software Programmer's Guide.*)

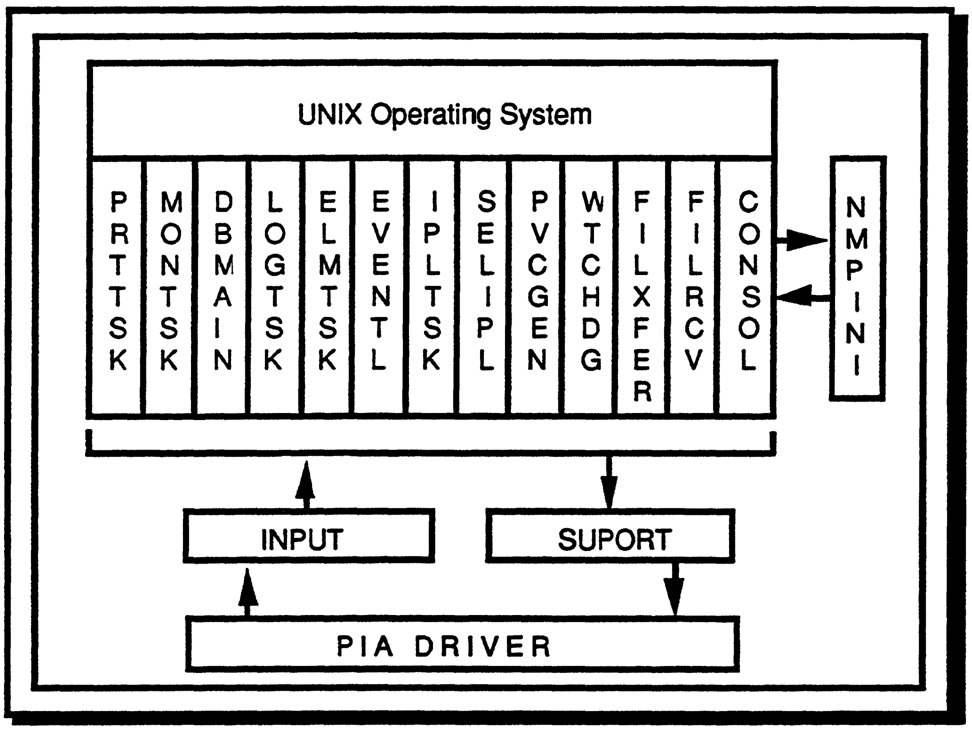

Figure 3-2 N7900 Software Overview

The three (3) tasks that are used almost exclusively on the N7900 Network Manager are as follows: *consol*, *input*, and *suport* tasks. These are the actual task names and the correct spelling of the task. The *consol* task is designed to support either the N7900s color display monitor or a VT100 terminal. It has several switch options available but the primary modes of operation are: first, to support the graphics capabilities of the Network Manager (through the use of *gcon* or *consol -g*); and second, to support a standard monochrome terminal via the *consol* task. The *consol* task is the operator's interface to the network and the database; it represents the sole interface between the node and the operator of the N7900. The *consol* program is menu-driven and responds to the single character commands shown in the menus. By using this task the operator can access, change, and administer all network resources. Refer to Figure 3-3.

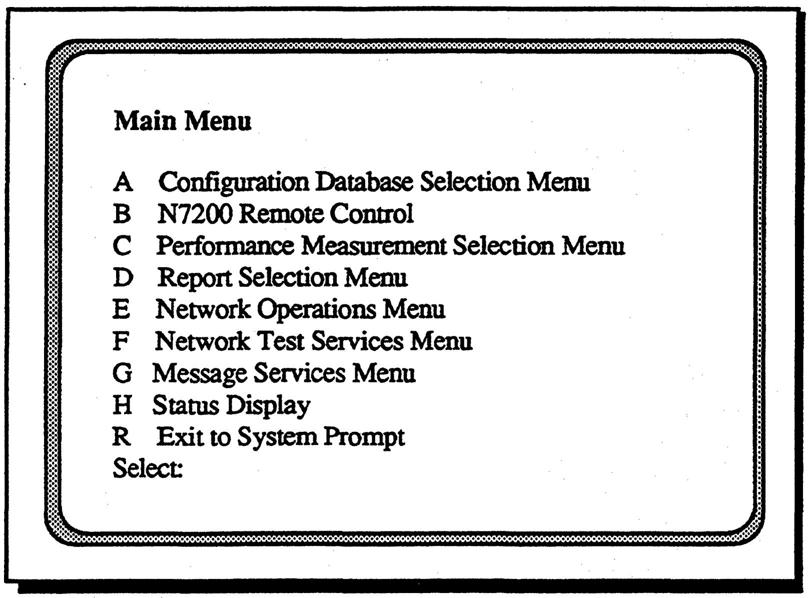

Figure 3-3 N7900 Main Menu Selection Screen

All activity and messages for input and output to the N7900 is accomplished through the *support* and *input* tasks. The *suport* task interfaces to the Packet Processor Module (PPM) through the parallel interface adapter. It is an integral part of the N7900s shared memory, and can read messages from other N7900 tasks via its input message queue. These messages are then written to the PIA interface. The *input* task receives all messages from the network via the PIA driver/interface and makes this information available to the other resident tasks in the N7900. Also, the *input* task determines where the message or data should be placed via a dispatch table. Once this is accomplished the task responsible for storing, printing, or logging the information is executed.

Database Management

The database management task (*dbmain*) is called from other tasks to perform the storage and retrieval of statistical records, configuration data, subscriber information, call requests, and all file and record update operations. For example, the *consol* task uses it to initially create the network database, updates, changes, and all configuration parameters. The *suport* task uses it to pass call request information for processing.

Communications with the database manager are performed using standard UNIX directives. A resident common area, in the N7900, referred to as BUFFER is used for all records. This provides an area where the records can be updated, read, and referenced by other tasks. Flags are used to inform the *dbmain* task when a specific operation is completed thereby freeing the database manager task. Associated with the *dbmain* task are many subroutines and other tasks being invoked to support the database operations. Any operation that alters the database is flagged and passed on to the database logging task (i.e., *logtsk*).

Some of the files supported by the database functions are:

- LP Configuration
- Port Configuration
- Node Configuration
- SVC Subscriber
- PVC Connection
- Group Send
- Group Name
- PAD Profile Parameters
- PAD Management
- System Information

Database Logging Task

The function of the database logging task (*logtsk*) is to keep track of the status of the nodes in the network, perform database recovery operations, transfer of mastership from one N7900 to another, distribution of database changes, and log all changes that have been made to the master database.

Once this routine is initialized it reads the node configuration file and establishes the node status table (NODTAB.DAT). This table is used to determine the attributes of the nodes in the network (i.e., is the node insync, is the node attached to an N7900, and is it operational). Then it opens the database log file (LOGFIL.DAT) or creates one if it does not exist. From here on in, it monitors all transactions, that is; node operational messages, node not reachable messages, request for mastership, and local database changes via the operator.

If the database manager task is successful in making the required change, it formats a message and sends it to the *logtsk*. The database logging task then adds the record to the database log file, sends a message to the event log task including a sequence number, node ID, and function performed. It then requests that all N7900s and N7500s in the network be automatically updated via the APPLY message to ensure database concurrency.

Event Log Task

The event log task (*eventl*) performs the following functions:

- Maintains a disk log file of events
- Processes requests from the console task for event log records
- Processes remote database requests for database operations and event log records

The event log task first attempts to open the event log file (ALMFIL.DAT). If the file cannot be opened, the program creates and initializes a new version of the file. Then *eventl* calls a routine to build an index table for accessing the event log records by time. Now *eventl* continuously loops receiving and processing messages until it receives an end message from the database manager at which point in time it closes the event log file and exits.

Requests for event log records can also come from the *consol* program. If the *consol* program desires to display or print an event log message, it can retrieve the records in three (3) ways: first, it can retrieve the last record in the file; second, retrieve the first record in the file that has a time stamp that is greater than or equal to a specified time; and third, retrieve the record previous to the last record retrieved.

Requests can come from the remote console task or from the Computer Interface Process (CIP). CIP handles requests from remote subscribers or computers to perform file and database operations or get event log records. If the operation involves changing a file or a record (i.e., add, update, or delete), the remote source must send a request to invoke maintenance to insure that the node is master and can make the requested changes. When a remote subscriber requests an event log record the message is received by *suport* from the CIP process and forwarded to the event log task (*eventl*).

PVC Generation Task

The PVC generation task (*pvcgen*) performs the automatic generation of PVCs between the current node and the other nodes in the network as part of the IPL process. The IPL task sends a *pvcgen* message to generate all of the PVCs in which both ends of the connection reside on the current node. Then, the PVC generator task is requested to generate any PVC connections between this node and any newly operational remote node as well as receiving, routing, and logging all PVCs created and removed from the Level III process in the PPM.

When the PVC generator receives a request to generate the PVCs from the current node to some remote node, it first checks the node status table (NODTAB.DAT) to see if they have been generated. The node status table contains a single byte per possible node and is indexed by the node ID. Each bit in the status byte has a meaning. The PVC generator uses bit 5 to indicate whether PVCs have already been generated to the node in question. PVCs are created in three places:

- As part of the IPL process
- The operator may elect to create or disconnect any PVC while the system is operational
- The operator may elect to create or delete a connection from a PVC subscriber to a test service

File Transfer Task

Just as the section title suggests this task informs the file transfer operation (*filxfer*). The transfer file option is reached via the network operator selection on the main menu or through the CIP process. Once the selection is made, the operator is prompted for information regarding the source and destination file names, path names, and the source and destination node identification.

This allows any non-database file to be transferred over the network from one node to another. There are two (2) tasks that are invoked automatically by the *suport* task to implement file transfer, that is, the transfer task (*xfrfil*) and the receive task (*rcvfil*). The transfer task actually sends the specified file over the network and the receive task performs the error checking, creates the destination file, and receives the messages from the transfer task.

Initial Program Load Task

The IPL task (*ipltsk*) performs the function of loading, and configuring the PPM as well as all LPs on the node. The majority of the operation of this task is actually controlled by the PPM with the N7900 operating in a slave environment. Details and an example of what operations are required is covered in Appendix D.

The basic operation is that the PPM requests the IPL to begin with an EDBCC message. This routine responds with an ADBCR message and then initializes the node status table by turning off all status bits except the one that indicates that the node is in the configuration file. Then the PPM requests the node configuration list with the *XIPLC* message and routine. This routine generates the node configuration table which is then sent to the PPM. First the FEP load image is sent via the *XCDT* routine followed by the PPM load image which uses the *XMDT* routine.

Once the downline load is complete, a DLLC message is sent from the PPM to the N7900 event log file to indicate that the downline load has been successfully accomplished or if there are any discrepancies. For example, a list of LPs that are in the database but does not exist. Now the PPM sends an IPL2 message to request the node and LP port parameters.

The following messages are sent from the N7900:

EDBCC	Start IPL sequence
XIPLC	IPL request for configuration table
XCDT	IPL request for block of LP load image
XMDT	IPL request for block of load image
DLLC	IPL downline load completion
IPL2	IPL start second phase IPL
IPL3	IPL start third phase IPL

The following message is received in the N7900:

ADBCR	Acknowledge IPL start

One of the last routines that is executed is used to generate the statistics parameters messages. This is through an IPL3 message. The last task is used to establish any PVCs for this node with the PVC generator task. A message is sent to generate the PVCs that are totally contained within this node; then the routine exists.

Remote IPL Task

This task is invoked to load and initialize an individual LP (*selipl*). This task may be initiated by either the operator at the N7900 or by the LP. If it is operator requested, the *consol* program sends a IPLREQ message to the PPM on the specified node. The PPM responds with a FPSLOK message which indicates whether the response has been accepted. When the node and LP are in a state to process, a FIPLOK message is sent to the N7900 and the *suport* task starts the selective load task.

If an LP is manually reset, it sends a message to the PPM indicating its LP ID. The PPM interprets the request for a selective load and as soon as the PPM software is ready for the selective load, a FEPREQ message is sent to the N7900 Network Manager. The receipt of this message by the *suport* task causes it to start the selective load task.

The selective load process consists of sending the appropriate load image, node and port parameters, and LP statistics parameters for any PVCs specifically connected to this LP.

Elect Master Task

This task (*elmtsk*) is invoked when a node requests master status for database operations and no master exists. When an operator requests master status and no master exists, the *consol* program asks whether to initiate the elect master procedure. The task then has every node exchange its latest file update IDs with all other nodes, and the lowest node ID with the latest file update ID becomes master.

Watch Dog Timer Task

The watch dog timer task (*wdttsk*) function is designed to monitor the health of the locally attached N7900 and also to notify the operator of incoming calls from remote nodes or subscribers. The first function of the watch dog timer task on the N7900 is to receive timer requests which are sent once each minute from the locally attached node. The *suport* task receives a WDTREQ message and expects a series of three (3) replies each with a different status message number 1, 2, and 3; respectively. The reply message is a (WDTRES) message with the appropriate status message. Each status message signifies which parts of the N7900 is operational and this allows the attached PPM to immediately detect any error conditions.

The second function handled by *wdttsk* is to notify the operator of any incoming calls from a remote node or subscriber. When an incoming call is received in the PPM, the message service process will display that a call is waiting, along with the subscriber number, node and the name.

Ancillary Tasks

Other tasks in the N7900 are fairly straightforward and need little or no explanation. The print task (*prttsk*) is responsible for display/printing events as they happen. The monitor status task (*montsk*) is responsible for monitoring the network status and the NMP initialization task (*nmpini*) is activated when starting the N7900 Network Manager or initializing other N7900 tasks.

Centralized vs. Distributed Control

Your well aware of the many capabilities and features associated with the N7900 Network Manager. Centralized versus distributed control of networks operations has been a concern for many corporations that use networks. The N7900 can operate in either mode with its master database which only allows one operator to update the database, after the proper password and security measures have been followed.

Some concerns stem from security issues; another area of concern is loss of control of the network administration functions. Some organizations want to manage the network and all of its resources from a central location. This would provide better control and tighter security. Other companies are looking at the entire network and base their decision on network availability, load sharing, and redundancy. This especially holds true when the network is large, spans time zones, or is composed of an array of off site or remote locations.

It is not our intention to tell you which one is better for your particular site or organization. The ultimate decision is up to you and your organization. Our intention as a networking company is basically to relate what our networking consultants, network troubleshooters, and designers use as a guideline. From this we anticipate you will have sufficient knowledge to comprehend your network performance requirements.

The following guidelines or recommendations were derived by our networking consultants in successfully solving networking problems. They are not in any specific order. Some of them do not have to deal directly with the N7900, per se, but overall system/network procedures that are proven to work.

First, when selecting line speed based on estimated traffic, try to keep the line utilization around 50%. This allows for traffic growth and helps carry traffic bursts gracefully. If a line is over 50% utilized, you can provide an additional path to reduce this utilization.

Second, build network topologies that provide diversity routing. This means that the dynamic router will have several choices available to maintain fast data delivery. Ring, loop, or mesh topologies are best. Automatic load balancing and higher throughput are provided by using parallel paths between high traffic sites.

Third, with low internodal operations (i.e., trunk speeds of 9600 bps or below) the number of nodes from the source to the destination should be kept low. More than two (2) nodes in series at this speed imposes noticeable delay due to transmission speeds. One means of improving this is to provide multiple paths to raise the available bandwidth.

Fourth, with internodal trunk speed in excess of 56 Kbps, the number of nodes in a series is less of a concern. A usable guideline is five (5) hops for the longest path in the network with 56 K or 64 Kbps internodal trunks.

Fifth, allocate one N7900 per 5 to 10 nodes. For networks that require high availability, improved network operations, or where a line failure could cause the network to become partitioned two (2) Network Managers are suggested or an N7900 and one N7500.

Sixth, the Network Manager and the attached node must be power cycled together since the PIA is in continuous operation. A power cycle on only one side of the interface can cause an indeterminate interrupt to occur. For high availability applications, an Uninterruptible Power Source (UPS) is highly recommended. This will eliminate network downtime from local power disturbances.

Seventh, the use of N7900s at remote sites will provide additional call processing, load sharing, and redundancy.

Eighth, select internodal trunk speeds at least as high as the highest subscriber port speed. If this is impossible, use multiple trunks to provide the required bandwidth.

Ninth, when multiple trunks are utilized in the network to provide alternate paths and congestion control, select different LPs (end points) to further avoid congestion inside the node.

Multi Level Security

A major concern to network administrators is security. In order to protect your network AMNET provides several distinct levels of password protection making the network secure. The password can be up to six (6) characters long and can include any alphanumeric character. Further, this password is case sensitive; i.e., it will recognize the difference between a uppercase letter 'A' versus a lowercase letter 'a'. The five (5) levels of security password protection are:

- Computer Interface Password
- Low-level Operations Password
- High-level Operations Password
- Login Password
- Maintenance Password

The High- and Low-level Operations passwords protect portions of the database and certain network operations. For example, you need an operations password to add, modify, and delete records in a specified database file. Also, as part of the *Change Password* facility you must have a High-level password in order to change passwords on the network. The High- and Low-level passwords are network wide passwords. The Computer Interface Process password provides protection from unauthorized subscribers access to the Computer Interface Process.

The Login and Maintenance passwords are used for security access into the N7900. In order to gain access to the N7900 the network operator must use the correct Login password . The last password is the Maintenance password which is used for support of the UNIX environment on the N7900. This is equivalent to the "root" password or "super user" password and is used to support the operations outside the realm of the N7900 software. The Login and Maintenance passwords can be different for each N7900 or N7500 in the network.

SECTION 4

NETWORK SERVICES

Network Services

In order to understand the capabilities of the network and the Nucleus 7400 and Nucleus 7500 this section provides you with the following information:

- X.25 Capabilities
- Virtual Circuits
- Multiple Trunk Protocols
- Network Routing
- Additional N7400 and N7500 Capabilities

The X.25 Connection

The NSNs support X.25 in accordance with the 1984 CCITT Recommendations. Both Switched Virtual Circuit (SVC) and Permanent Virtual Circuit (PVC) services are supported. A wide range of facilities, known as X.2 facilities, are supported to enhance network services, performance, and control. The following X.25 facilities are supported on a per port basis:

- Lowest Incoming Channel
- Highest Incoming Channel
- Lowest Two-way Channel
- Highest Two-way Channel
- Lowest Outgoing Channel
- Highest Outgoing Channel
- Default Packet Window
- Default Packet Size
- Maximum Packet Size
- Flow Control Negotiation
- Throughput Class Negotiation
- RPOA Selection
- Default Calling Address
- Local Charging Prevention
- Transit Delay Selection and Indication

This level of flexibility provides an important tool to network designers and administrators for optimizing system performance. The process of integrating the network with other X.25 devices is also simplified by having control of these key interface characteristics. To further enhance these capabilities the following line interfaces are available on both the N7400 and N7500:

- RS-232-C
- V.24/V.28
- RS-530
- V.35

In addition to the DTE or DCE logical appearance setting, the NSNs allow you to coordinate the trunk parameters to meet the characteristics of the X.25 facilities being used and to achieve a better match with different networks and/or X.25 based devices. Table 4-1 shows some of the X.25 default parameters stored in the database. These values can be altered to meet your specific requirements through a simple menu selection mechanism.

Parameter	Valid Values	Default Value
SVC LCN	1 - 4095	10
Incoming Address Check	Y/N	N
DCE/DTE Restart Timeout (sec)	5 - 360	60
DCE/DTE Incoming Call Timeout (sec)	5 - 360	180
DCE/DTE Reset Timeout (sec)	5 - 360	60
DCE/DTE Clear Timeout (sec)	5 - 360	60
Default Packet Size (bytes)	16 - 4096	128
Default Window Size (packets)	1 - 7	2
DTE Clear Retransmit Count	0 - 10	0
DTE Reset Retransmit Count	0 - 10	0
DTE Restart Retransmit Count	0 - 10	0

Table 4-1 X.25 Parameters

Virtual Circuits

The most comprehensive feature of packet switching is its ability to have virtual circuits, also known as logical channels. This means that it is possible to support multiple, simultaneous, communications sessions with only one physical connection - the X.25 trunk.

It is also possible to select the number and type of virtual circuits (switched or permanent) to achieve optimum use of the X.25 trunk. The number of virtual circuits dictates the number of simultaneous sessions that can be supported over the X.25 trunk. The N7400 supports 264 simultaneous virtual connections and the N7500 supports up to 2,000 virtual connections.

Switched Virtual Circuit

A Switched Virtual Circuit (SVC) is the virtual connection initiated by a subscriber. A virtual call is set up dynamically when one device attached to the NSN needs to communicate with another device. The SVC is identified by a unique address or subscriber ID. Within that subscriber record is the subscriber ID, subscriber name, port number, LCN, and Hunt Group if appropriate. As part of the call setup process, call information included in the call request is validated against information maintained in the subscriber file in the database. Subscribers can also elect to have their calls redirected on two conditions; that is, busy or out of order. This call redirection can be established to support up to five (5) subscriber IDs. As part of the call establishment the following services are provided:

- Subscriber Address Confirmation (calling and called)
- Facilities Validation
- Address Mapping (logical to physical)
- Establishment of Network Connection
- Call Barred
    - Incoming Calls
    - Outgoing Calls

SVC Subscriber Addresses

A subscriber ID is assigned to a port as an identifier for a subscriber who will use that port. Calls are routed in the NSNs based on the contents of the called address and user data fields of the Call Request packets. Subscriber IDs are used to logically identify users. The use of wildcards in a subscriber ID eliminates the need to specify the ID of every subscriber who will use that port. The SVC addressing format can be set up according to the needs of the individual network and conforms to CCITT X.121.

The following characters are allowable as part of the Subscribers ID:

- Digits 0 - 9
- Wildcard characters are allowed:
    - * matches any group of digits
    - ? is a single digit match

The '?' means that a digit must be present, but it can be any digit. The '*' matches any sequence of digits, including no digits. For example, 3110* is any address that starts with 3110, including 3110 alone. The address *00 is any address that ends in 00. And the combination of 3110*00 is any address that starts with 3110 and ends with 00.

Abbreviated Address

For Async PAD subscribers an Abbreviated Address or mnemonic address (alias) can be defined to describe a subscriber. For the Abbreviated Address any character is valid except the '+' sign. The operator can establish the alias table for the Subscriber ID with one (1) to ten (10) characters. It is recommended that the first character be a letter to avoid confusing this address with an actual subscriber ID.

Permanent Virtual Circuit

Permanent Virtual Circuit are defined as a part of the subscriber configuration process. A PVC is a permanent logical connection between two end points over the network. It is comparable to a point-to-point leased line. When you select a PVC you require that a device communicates with a pre-defined destination. The use of the PVC eliminates the need for the user to initiate any connect or disconnect procedures, and no call setup or call clearing is required. In order to use PVCs the following must be defined:

- Source and destination Subscriber IDs
- The logical channel numbers that apply to each end of the PVC
- Parameters that apply to the PVC connection; i.e., PVC ID

PVC connections are then created in the network, either automatically during the node's Initial Program Load (IPL), or under operator control while the network is on-line. When a PVC is created in the network, all nodes establish a table entry that describes the two end points of the PVC connection. This is the only pre-allocation of resources that occurs. When data is applied to the designated logical channel, it is automatically routed by the network (using normal internal routing procedures) to the other end of the connection. Wildcard characters are not allowed in the establishment of PVC IDs.

Additional User Facilities

The following CCITT Facilities are implemented with the N7400 and N7500:

- Hunt Group
- Group Send
- Closed User Group (CUG)
- Fast Select Accept
- Call Redirection
- Reverse Charging Accept

For facilities that apply to SVC service, a subscriber may negotiate use on a per call basis. As part of the database configuration process, some facility usage rights (e.g., Reverse Charging Accept, Group Send) are defined for each subscriber. These subscriber records also include any qualifying information that is required as the CUG name and maximum priority level allowed. The database is then accessed during a call request for a facility request validation. If valid, the facilities are used for the duration of the call.

Characteristics that apply for PVCs are defined as a permanent parameter in the database. As PVCs are created, these are loaded into the node and automatically applied.

Facility usage rights can be changed while the network is on-line. For SVC service the database change is automatically effective when the next call is established. For PVC service, the PVC is recreated on-line by the operator, causing the new characteristics to take effect.

Hunt Group

The Hunt Group feature allows numerous X.25 ports (up to 32) to be defined in the N7900 database as a single subscriber SVC address. The network operator has the flexibility to establish this resource under a single address rather than having several different addresses for a single location. This eliminates the problem caused by a port failure or network congestion with over 250 alternate addresses and supporting three (3) algorithms to further enhance the Hunt Group facility. The available algorithms are:

- Round Robin
- Alternate Route On-Line Busy
- Alternate Route On-Line Failure

When the *Round Robin* algorithm is used the first call to the Hunt Group goes to the first port. The second to the second, and so on, going back to the first port after it has reached the last. If the selected line is not up, the next available line is selected. This allows the load to be distributed over multiple ports and provides back-up in the case of a failed port.

When *Alternate Route On-Line Busy* is chosen the first line that is not busy (LCNs available) is selected for the call. If the selected line is not up, the next available line is selected.

When *Alternate Route On-Line Failure* is used the first line that is not failed is selected as the line for the call. If this line is busy then the call is cancelled even if other lines are available.

Group Send

Group Send is a service that allows a subscriber to broadcast a message to a pre-defined group of subscribers through a single call request. To use this service, the originating subscriber places a call to the Group Name established by the operator. The call request packet is established with the Group Send ID as the calling address with the Group Send Identifier of 0901 in the facility field.

The call request is sent in parallel to all members of the group. When all subscribers have sent either a call accept or a call clear, the message is sent. The number of subscribers that can be assigned to a group is equal to the number of simultaneous connections available on AMNET's NSNs. For the N7400 this value is 264, and for the N7500 is 2,000. The number of groups that can be defined is up to 10,000 maximum.

Closed User Group

Closed User Group (CUG) is a privacy feature that allows creation of groups of subscribers [up to ten (10) CUGs per subscriber] within the network. Members of a CUG can communicate with other members but not with network users outside of the CUG. Closed User Groups can have a value from 01 through 99. Since a subscriber can belong to multiple groups, we establish an index in the CUG table for the subscribers selection.

When the call request is placed to a CUG an '03' followed by the two (2) digit CUG number in the facility field is used. If this is not established, then the call will be placed to the default or Preferential CUG. All of the test services (data drop, data looping, and message generation) are members of all Closed User Groups and have incoming access. Message services is a member of all CUGs and has both incoming and outgoing access.

Fast Select Accept

The Fast Select Accept is an optional X.25 facility that authorizes the NSNs to transmit incoming calls that request the X.25 Fast Select facility. This is on a per subscriber basis and if it is set to 'No' then the network will not send the call request packets with the Fast Select facility to the subscriber.

Call Redirection

Call Redirection is a network service that allows a subscriber to be dynamically reassigned to an alternate location. There are three (3) available redirections: Systematic Redirection, Redirect on Out of Order, and Redirect on Busy. Systematic Redirection automatically redirected all calls to the alternate subscriber specified in the first redirected subscriber number established by the operator. Up to five (5) alternate subscribers can be established. Calls will be attempted in consecutive order to this specified list in the order that the list was established.

Redirect on Out of Order is similar to Systematic Redirection except that the call is redirected only if the subscriber device is out of order. Redirect on Busy is analogous to Redirect on Out of Order except that the call is redirected only if the subscriber device is busy.

Reverse Charging Acceptance

Reverse Charging Acceptance is an optional X.25 facility based on a per subscriber basis. Typically, subscribers are allowed to send call requests with Reverse Charging to the network. If this parameter is set to 'No', then the network will not send call request packets with this facility to the subscriber.

Flow Control Negotiation

Flow Control Negotiation is an optional X.25 facility that permits negotiation of the Flow Control parameters that is the packet and window sizes for each direction of data transmission on a per port basis. When selected all call requests and call accept packets sent by the network will contain packet size and window size negotiation.

Throughput Class Negotiation

Throughput Class Negotiation is an optional X.25 facility that is specified on a per port basis and permits negotiation of the Throughput Classes for each direction of data transmission on a per call basis. Throughput Class Negotiation may be included in call request and call accept packets.

Network User ID

This is an optional X.25 facility that enables the subscriber to provide information to the network on a per subscriber basis for billing security or network management. The network operator can elect to have the Network User ID (NUI) required by all subscribers as optional or disabled. If the calling subscriber is attached to a logical DTE port the value of this parameter is overridden and NUI will not be allowed. When NUI is included in the call request packet it will be added to the call accounting file. If the destination port for the call is a DTE the NUI will be forwarded to the port and will be included in the outgoing call request packet. If the destination port is a DCE the NUI will not be present in the outgoing call request packet.

Network Services Addressing

The N7400 and N7500 Network Switches reserve certain addresses for special network services. This allows a subscriber or operator to access test services and other special features in the network by using standard X.25 Call Request procedures. These test services are covered in greater detail in Section 5.

Some of the test services provided are:

- Data Drop - data sent to this address is discarded
- Data Looping - data sent to this address is looped back to the source
- Computer Interface Process - for connecting to the Network Manager Console operations

Call Accounting

The Call Accounting feature provides network accounting data for SVC calls and PVCs based on the following reported events - Call Setup, Interim Report, and Call Clear. Totals of packets and bytes sent and received for each connection are collected by the node and written to a Call Accounting file on a designated Network Manager. The Call Accounting file is then available for billing report generation or transfer to another system for data analysis.

The *Call Setup* is generated when the SVC (call) or PVC is established, this record contains a time stamp and all of the data known to the N7900 at setup time, including subscriber addresses, end nodes, LP and port, throughput class, and facilities requested.

The *Interim Report* is sent at configurable intervals between the Call Setup and the Call Clear. This interval is configurable in six (6) minute increments from six (6) minutes to 25.5 hours. This record contains the call identification, time stamp, as well as the number of packets, segments, and bytes sent and received since the last Interim Report or Call Setup.

The *Call Clear* is produced when the SVC call is cleared or the PVC is deleted. This record contains the call identification, time stamp, reason for termination (clearing cause and diagnostic codes), and counts of packets and bytes since the last Interim Report or Call Setup.

Operator Functions

The operator can request Call Accounting file rollover. This will cause the Call Accounting file to be closed and a new file to be initialized. For example, if you want to have a separate Call Accounting file for each day, you could use the file rollover operation to automate this process. This is automatically accomplished by the N7900 Network Manager once Call Accounting function is enabled. Monitoring the disk free space is left to the operator. The files created and stored on the system disk could be large, depending on the network traffic, and the possibility of running out of disk space is present. The file may be transferred to floppy disk and processed by a user supplied utility program.

Packet Efficiency

The amount of data in a packet is variable and has an effect on the response time and charges incurred over the network. Usually transmitting full packets is the most cost-effective, but it may not be the ideal answer to your application.

There is an efficiency break point in the packet size. In order to understand this phenomena of packet efficiency, Table 4-2 illustrates the number of packets per second at different link speeds versus packet size. It should be obvious by looking at the difference between a 256 byte packet and a 1,024 byte packet operating on a 56 Kbps trunk.

PACKET SIZE	BITS/ FRAME	PPS @ Line Speed 56 Kbps	PPS @ Line Speed 14.4 Kbps	PPS @ Line Speed 9.6 Kbps
1024	8264	6.7	1.75	1.2
256	2128	26	6.7	4.5
128	1104	51	13	9
64	592	95	24	16
32	336	166	43	29
16	208	269	69	46

Table 4-2 Packet/Trunk Efficiency

SECTION 5

NETWORK OPERATIONS

Status & Control

Control and management of the NSNs are supported by the N7900. The Network Manager provides operator access to database management, monitoring ability, connection information, and maintenance/diagnostic functions. This information is available on-line and in real-time so that any changes in network topology, congestion, and excessive line errors are reported directly to the operator.

The N7900 has a command structure based on user-friendly interactive menus, on-line 'Help' facility, pop-up menus and prompts. Prompts are used to indicate the types of responses expected from the operator and, where applicable, to display a list of valid responses. Menus are used primarily to display command options to the operator. Built-in checks and safety functions are included to reduce operator error.

The N7900 provides detailed on-line status for every component in the network. The network operator can customize alarms and alarm thresholds to tailor the network reporting for your particular environment. Further, the network operator has the capability of: controlling all network components; providing on-line functionality such as taking lines or individual VCs in and out of service while the node is on-line and not having to re-IPL the node; restarting hardware components; performance testing; gather status information including utilization; and error rates on a real-time basis.

The real-time information generated by the nodes can also be collected for off-line analysis. Information is collected by the individual nodes and transmitted to the N7900 site at pre-defined intervals or on a demand basis. This provides for a flexible implementation while minimizing the loading that data collection places on the network.

Generally, the types of data logged are alarm indications and event reports. Alarms on an individual component or group of components can be recorded in a file, and/or logged to the printer. This recording of events allows the network operator to closely monitor network and component performance. Alarm information can be directed to specific locations throughout the network.

In order to better understand the capabilities of the N7900 this section covers what an operator can establish, change and view from the Network Manager. It also covers the statistics and reports that are available from the Network Manager (N7900) on a real-time basis.

On-Line Control and Operations

On-Line Control is performed while the network is up and operational. It provides extensive configuration capabilities available to the network operator. Menus allow you to: enable or disable specific ports; examine or change port parameters; create or delete PVCs; transfer files for software maintenance and upgrades; check database concurrency operations; change configuration parameters; remove an N7900 from service for diagnostic or preventative maintenance; and view all events and alarms throughout the network.

These on-line changes take effect immediately but, for the most part, are only temporary changes. They stay in effect as long as the node, LP, or port that received the change(s) is not rebooted or re-IPLed. If the changes prove to be beneficial to network performance or network integrity the operator can make the changes permanent. Figure 5-1 illustrates this screen for the network operator.

Network Operations Menu

A On-line Port Operations
B NMP/Node Operations
C Database Concurrency Operations
D On-line PVC Create/Delete
E Change Level of Event Printing
F File Transfer
G On-line Update of Network Defaults
R Return to Main Menu
Select:

Figure 5-1 Network Operations Menu

On-Line Port Operations

When you select the On-Line Port operation, you have total control of any port or trunk in the entire network. Figure 5-2 illustrates the On-Line Port operations. These operations are only performed while the network is operational and will cause alarms and events to be generated at the N7900 since you have the potential to bring down an active port or trunk.

On-line Port Operations Menu

A Port Parameters Examine/Change
B Enable Port
C Disable Port
D Async PAD Profile Examine/Change
E Enable SNA Port Autocall
F Enable Logical Trunk Channel
G Disable Logical Trunk Channel
R Return to Previous Menu
Select:

Figure 5-2 On-Line Operations Menu

Port parameters can be examined and changed on-line but if any calls are active on that port they will be cleared. This allows the operator to view the current setting of any port in the network since the last IPL. The values you see are read directly from the specified LP and not from the database therefore any prior on-line changes are reflected. Again, most of the changes are on a temporarily basis.

Depending on the port configuration some of the status messages that might appear at the workstation would be as follows:

- Link Disabled
- Link Failed
- Link Disconnected
- Link Up

To disable a port you must specify the node ID, LP number and port number. While the port is disabled the port status is reported as 'Link Disabled' on the status display. In this mode, the port cannot send or receive packets. To reinitialize or enable a port, you must specify the node ID, LP number and port number.

When the Async PAD Profile option is selected you can modify any parameter of the integrated Async PAD in the network. Making changes to the PAD profile will not impact the data traffic on that port. You also have the option of making the changes permanent or keeping the PAD profile; that is, keeping the old profile parameters in effect.

During critical subscriber operations or maybe even for downloading the latest versions of the operating system to the nodes, the operator may want some additional assurance as to availability of trunk lines in the network. Instead of going through a reconfiguration session for a particular node the alternative is to allocate additional back-up trunk lines. This is accomplished via the use of the Enable Logical Trunk Channel. This operation allows the operator to enable a Logical Trunk for a port that already has been defined in the database. To cancel this operation there is the Disable Logical Trunk Channel.

N7900/Node Operations

Another component of the N7900 is its ability to handle Call Accounting information and the ability to automatically rollover this information to the systems hard disk. The operator is allowed to re-IPL (restart) any node in the network. This is useful in the event that the node is experiencing severe problems or the operator wants to download a new software release to the node. Another useful feature available to the operator is the selection of an alternate network manager. With distributed N7900s established in the network, the operator can select an alternate network manager when the designated back-up N7900 is taken out of service. This could be for a software upgrade, or for preventative maintenance.

Accounting Information

The Call Accounting feature provides network accounting data for SVC and PVCs calls. Totals of packets and bytes sent and received for each connection are collected by the node and written to a Call Accounting file on a designated N7900. The Call Accounting file is then available for billing report generation or transfer to another system for data analysis.

For each connection, the following information is available: (1) a time stamp of when the connection was established including subscriber IDs, end nodes, and facilities requested; (2) an automatic interim report [selectable on six (6) minute increments] containing the call identification, byte count, and the number of packets since the connection was established or since the last time this interim report was available; and (3) a time stamp of when the connection was cleared including subscriber IDs, reason for termination, call identification, and a count of the number of packets and bytes since the connection was established or since the last time the interim report was run.

The automatic operation for rollover is useful when you want an account on a daily basis for the aforementioned statistics or if you want to transfer this information to a floppy disk for further off-line processing. In order for this information to be available and valid, the Call Accounting function has to be enabled. Once the Call Accounting is enabled it will collect the data once the system is operational.

There are safeguards built into the system. Each safeguard is accompanied by an alarm message generated to the N7900. The primary reason is that it is possible to physically run out of disk space on the system disk by collecting the Call Accounting information and not processing it properly. First, an alarm message is sent to the N7900 relating that the disk space is diminishing to an unacceptable level. At this level Call Accounting will not accept any new calls for monitoring.

If the problem continues and the operator has not performed any corrective action the second level will be the complete shutdown of the N7900. That is, all the N7900 processes are stopped and the operator must transfer or delete the Call Accounting file from the system disk and then restart the Network Manager. These are some of the automatic safeguards built into the N7900.

There is an AMNET utility called *'caf'* which allows the operator to monitor, manage, and selectively view records contained in the Call Accounting file. The operator can view, on the attached monitor, the Call Accounting file based on the following criteria: look for all records of a local subscriber ID or a remote subscriber ID; view records based on the LCNs whether it is a local LCN or a remote LCN; select records based on a certain reason codes; or select all records. Some of the selectable reason codes are:

- PVC Call Setup
- Restart Request
- SVC Call Request
- SVC Cleared Call
- Line Down
- PVC Removed

Further, the operator can transfer this file to a floppy disk for additional processing, print the file to the printer, or display the entire file by selecting all records.

Change Passwords

In order to maintain security on the network and to protect the database from unauthorized access, the network administrator should periodically change passwords. The operator can change the three (3) levels of password security available; i.e., low-level, high-level, and computer interface. All three (3) passwords are in the network database and therefore any change to any password is automatically disseminated to all nodes.

Boot Disk Operations

There are two additional operations that can be performed in supporting the NSNs. The first is to allow the operator to configure a floppy diskette for a specific node ID. You can create a bootable diskette for the node as well as create a complete diskette with all the software and configuration files. Now the floppy can be inserted into the specific node and the node can be powered on or re-IPLed. This is considered a local operation because it is performed at the N7900 console.

The second operation that can be performed is to actually download the software directly to the node. You can selectively download to the NSN the port configuration file, the LP configuration file, or the node configuration file. Furthermore, you can download the entire software update to the node. As the files are downloaded, their names will be displayed on the monitor and when the operation is complete an information message will be generated.

Database Concurrency Operations

Normally, only one N7900 has the status of being master for the entire network which prevents multiple operators from making simultaneous changes to the database. This is controlled through the use of a master status, which is passed around the network to other N7900s on request (similar in design to passing the token in a token ring network). Each time there is an update to the database, the update is automatically sent to all other N7900s and N7500s in the network, so that the data remains consistent and in synchronization.

Sometimes it is necessary to take an N7900 off-line. This could occur for preventative maintenance, software update, or when a network becomes physically partitioned. When one of these conditions occur it is possible, that the current N7900 database could be altered. When the original N7900 is re-introduced into the network, it may have to be brought back into synchronization with the rest of the network. In order to accomplish this we can execute a database recovery.

Database Recovery

The operator will automatically be notified if this problem exists when the N7900 is returned to an operational state. Like all of the other operations performed by the N7900 this operation is available through a simple menu selection; that is, *Database Recovery*. The current master node will actually execute the database recovery; since this node has the most up-to-date database; and you will be informed when the operation starts and when the database recovery has completed.

What happens if the network became physically partitioned? This can occur when we have multiple N7900s supporting a network and a link between two of the nodes becomes disconnected causing the network to split. Database changes were made to each N7900 since they are now supporting separate networks. When the link is re-established the N7900s database is not in synchronization and the operator will be informed that there are multiple database masters.

A determination has to be made as to which N7900 should continue to function as the master and which one should be cleared. In this situation the local operators would decide as to which N7900 should retain mastership of the network. The other N7900 is cleared of its remote mastership via a *Clear Remote Master* operation. The *Clear Remote Master* operation removes master status from the specified N7900 and automatically brings its database into sync with the current master N7900.

Now in large networks there could be several N7900s throughout the network. What about their status? The *Enable Backout* operation will bring into sync any other N7900s that were partitioned along with the previously cleared master. Again this information is automatically made available to the operator via a printed message on the master N7900s event logging printer.

On-line PVC Create/Delete

This allows the network operator to create a temporary PVC connection between existing subscribers (i.e., subscribers already established in the network database). You can provide a priority code when creating this PVC subscriber as well as cancel or delete this connection once it has been established. Messages will be generated, sent to the console and printed to verify that the PVC was created or removed.

This operation of creating a PVC subscriber is useful especially if several subscribers are experiencing problems. Various test services are available to check the viability of the line, connection, and to generate traffic or check load conditions for subscribers. The test services available are covered in the Network Test Services section of this document.

Change Level of Event Printing

All alarms and events are reported to the N7900. The N7900 includes a monitor, printer, and hard disk. With large networks and up to four (4) levels of severity information available the amount of information sent to the printer could be substantial. Even though we call them events, some of the events are informational only and do not require any operator action to be taken.

Some operators prefer to control the information logged to the printer. They may choose to have only Level I and Level II alarms logged. Both of these levels require operator intervention. To accomplish this we provide a means where the operator can select the severity level(s) to be printed. The operator can select to have printed:

- Level I only
- Level I and Level II
- Level I, Level II, and Level III
- Print all Levels

File Transfer

The *File Transfer* service allows an operator to initiate a file transfer between any two N7900s. This can be used for both system files (program update) as well as data files. The file transfer is designed to be used for software maintenance. Since files up to 256 Kbytes can be transferred some considerations and precautions should be given to the users on the network. In addition to specifying the file name you can specify a pathname where the file should reside. Files can be sent remotely but not deleted. Therefore, periodic file maintenance should be performed to delete old files.

Network Test Services

Network test services are provided so that an operator can:

- Verify that a subscriber is connected to the network
- Determine the integrity of the connection
- Check the performance levels of individual nodes and the network

These test services are primarily designed to check overall network integrity. But you can also verify PVC subscribers operations by using data drop or data looping. SVC subscribers can also use some of the capability provided by the network test services. An additional feature when using test services is that once you have initiated any test, you can exit from this routine and perform other network operations.

Some of the test services available are:

- PVC Data Looping
- PVC Data Drop
- PVC Message Generation
- Packet Trace

- Packet Delay Measurement
- Load Generator

PVC and SVC Data Test Operations

PVC Data Looping connects a PVC subscriber to a node, sends data to the node, then the destination node sends or loops the data back to the subscriber. The *PVC Data Drop* Service provides a network address that operates as a data sink which receives, acknowledges, and then destroys the packets. Both of these operations can be used by PVC subscribers who wish to verify that a connection to the network exists. The packet size used is the default packet size established for the network. This allows us to duplicate actual conditions on the network with the test services.

SVC subscribers access the Data Looping and Data Drop Service by placing a call (Call Request) with a called address of *999999nnn2*. The *nnn* is the test services node ID in the network. To run Data Drop for an SVC subscriber the Call Request is identical except the last digit is a 1; that is *999999nnn1*.

Any NSN on the network that has a monitor and keyboard provides you with an additional support capability - the Line Monitor Mode. This is especially useful when the operator selects the *PVC Message Generation* operation. Any operator on the destination node can directly view on the local monitor the messages as they comes across the network to this node. It operates similar to a data scope without any storage or playback capability.

Load Generation Operation

The *Load Generation* operation provides a tool for loading the network paths between two nodes with data packets. Several load generators may operate simultaneously. The operator may then measure performance characteristics of the network by using the other test services as *Packet Trace* and *Packet Delay Measurement*. As part of the *Load Generation* service the system will also inquire if you want to perform a *Packet Trace* as well as a *Packet Delay Measurement*, along with the *Load Generation.*

The *Packet Trace* service traces the path that a packet takes through the network. The messages printed lists all intermediate nodes each packet passes through as it travels to its final destination. Also, it keeps track of the number of hops required to the destination as well as the current node ID.

The *Packet Delay Measurement* indicates the time delay in milliseconds encountered for a packet to travel from the source node to a destination node and back again to the source node. This information is also available on the printer.

Reports

The Network Manager provides a variety of reports that describe the various elements of the system and its nodes. All of the generated reports will go directly to the printer and will not be displayed on the monitor. A summary of the available reports follows:

Configuration Summary Report

The Configuration Summary Report provides a concise inventory of the network components as defined in the database for every node in the network. The Configuration Report can be a lengthy document and, while this report is being printed, any system messages which would normally go to the printer will be temporarily held until the report has finished printing.

The report is generated by reading the Port Configuration File characteristics. It has a cross reference to it and is organized by the Node Configuration File, the Link Processor Configuration File, the Subscriber File, and the PVC Connection File.

Any inconsistencies in the network configuration will also be printed. These configuration parameter errors are reported to the network operator when initially entered. The operator may choose to ignore them but this report will generate a list of warnings and fatal errors encountered. Some of the possible errors are:

- Unused port
- Throughput class mismatch
- SVC assignment to a LP with no SVC support
- Undefined node
- Undefined LP
- Undefined LP type
- Incorrect LP type for a PVC

Subscriber Reports

The PVC Subscriber Report is a network wide inventory of PVC subscribers and is sequenced by subscriber ID. This report can be selected for a specific node or for all nodes. A summary is provided showing the total number of PVC subscribers, where each subscriber is connected to (source to destination address), and the LCN associated with that connection in the network.

The SVC Subscriber Report is a network wide inventory of SVC subscribers and is sequenced by subscriber ID. A summary is provided showing the total number of SVC subscribers in the network. The node, LP, and port number for the subscribers are also reported as well as any aliases associated with the SVCs.

Group Send Reports

The Group Send Report is really two reports in one. The first part provides a list of the subscribers that belong to each defined group name. The rest of the report is a cross reference report that lists all the members followed by their group name. All of the information comes from the Group Send File and and the Group Name File.

Alarms and Events

Alarms are the real-time status messages that appear on the printer ('Event Logging Printer'), archived to the system disk, and can be displayed via color-coded text ('Event Screen') on the Network Manager console. These alarms are automatically captured by the master N7900 and can be forwarded to two (2) additional Network Managers. This forwarding process is automatic and is usually implemented when the network spans time-zones, states, or countries.

Since the alarms, are automatically recorded to the system disk, a history of alarm activity can be developed for management reporting or for the development of trends. Furthermore, all alarms are time stamped and dated. Reports are presented with the most recent occurrence first, followed by any other alarms or events. Alarm reports are also customized by the operator. You can specify the severity level, time interval (start time and end time including the day, month, and year), a specific node or all nodes in network, and the name of the file to store the data on the system disk. You have the ability to monitor newly installed nodes on the network and to spot potential problems.

Also available is an *On-line Routing Maintenance* feature which allows you to specify other destinations for these reports. This is a simple mechanism to share the network responsibilities. Regardless of whether or not the alarms and events are displayed, they are always logged to the system disk. The operator may then retrieve the information at a later time to generate the report.

Four (4) levels of alarms are standard on the N7900. The alarm severity level is associated with a given message but it can be customized by the network administrator. The alarm levels and their meanings follows:

- Level 1 is a serious fault.

  This is a fatal error and immediate operator intervention is required. On the Event Screen this will be highlighted in red.

- Level 2 is a major fault.

  Action should be taken as soon as possible. On the Event Screen this major fault will be displayed in magenta

- Level 3 is a minor fault or traffic disturbance.

  The fault should be cleared at the first suitable occasion. On the Event Screen this fault will be reflected in yellow.

- Level 4 is an event notification.

  No action is usually required for most Level 4 events. On the Event Screen this level will be indicated in green.

Some of the error messages will appear with the letter 'S' or 'C' following the actual message number to indicate an actual status change. The letters are used primarily with messages whose text does not make the fault condition obvious. For example, we will use the letter 'S' to indicate that the condition has been 'Set' or entered and, conversely, the letter 'C' to indicated that the error condition has been 'Cleared' or exited. Thus, the status change could be generated as a result of the following:

- Trunk line up or down
- Port enabled or disabled
- Node unreachable (all trunks to that node are down)
- Node operational (node becomes reachable)
- Test service initiated or discontinued

With the color graphics interface the entire network topology is displayed at the N7900. The graphic interface is utilized in all modes of operations including configuration, monitoring, and database management. It provides you with a global view of your network's "health". You can view the network, subnetworks, individual nodes, and even display the EIA leads on the physical ports.

When an alarm is detected the icons representing the N7400, N7500, subnetwork(s), or N7900 will reflect the altered condition via a corresponding change in color. Also, any problems detected with the trunks lines will be reflected by a related change in color.

Statistics and Statistical Reports

The N7900 Network Managers' statistics package includes a comprehensive list for traffic, usage, and performance level measurements. These statistics are designed to provide a range of tools addressing the information requirements for your network installation. Statistics allow you to produce reports based on selected parameters and set statistics or threshold parameters. You need to specify the node, LP or port level statistics. For example, statistics may be used to:

- Monitor the utilization of critical system resources
- Isolate reported error conditions
- Provide actual performance counts
- Monitor for maximum values

The basic approach allows you to specify a series of parameters in the configuration database. These parameters are automatically loaded into the appropriate nodes and Line Processor cards. The maintenance of the 'Default Statistics Parameters File' (DSTPAR) in your database is entirely up to you. Once you have defined and verified the initial values you can then alter these values depending on the resource utilization at a specific node or nodes. This provides you with a simple scheme to set the monitoring thresholds and generate meaningful reports based on these values.

With the *Reports* option you can generate statistical summaries for all levels of the system or separate reports based on node, LP, or port level statistics. A summary of the available statistics follows:

Node Level Statistics

- Packets per second (data packets only)
- Bytes per second
- Congested Calls per second
- Call Requests per second
- Connections allocated per second
- Connections deleted per second
- Percent of queue elements in use
- Percent of packet processor run time
- Number of connections in use
- Percent of common memory buffers in use
- Percent of packet buffers in use

LP Level Statistics

- Percentage of buffers in use
- Number of queued messages
- Frames received
- Frames transmitted
- Frames rejected
- Frames retransmitted
- Percent of LP run time

Port Level Statistics

- Received frames per second
- Transmitted frames per second
- Retransmitted frames per second
- Frame Check Sequence (FCS) errors per second
- Reject frames issued (logical) for X.25 and trunk ports per second
- Ratio of retransmitted frames to successfully transmitted frames
- Ratio of rejected frames to accepted frames
- Percentage of port utilization

Statistical Threshold

If the measured value of a statistic goes above the default (operator specified) threshold, an alarm is generated. When the value returns to normal for a set period of time, another event is generated to inform the operator of a normal condition. Threshold reporting can be used to isolate specific error conditions or to indicate system resource shortages.

Computer Interface Process

The N7900 in conjunction with the NSNs incorporates a powerful service referred to as the Computer Interface Process (CIP). This CIP process permits a remote X.25 subscriber to access a subset of functions usually reserved for the network operator. This process requires access to the N7900 or to an N7500 that has ICP installed.

Through the use of CIP a majority of the functions, excluding graphics, are readily available to the remote subscriber. For example, access to node statistics, port parameters, and database operation as well as re-booting a node are all available to a remote subscriber.

CIP is only a member of the Closed User Group 99 and does not have either incoming or outgoing access. Therefore, any subscriber must be a member of Closed User Group 99 to place a call to CIP.

SVC subscribers can access the CIP process by placing a call (Call Request) with a special called address of *999999nnn9*. The *nnn* is the node number of the N7900 you wish to access. The remote computer subscriber places the call to the special test service address, enters the password in the user data field to allow the subscriber to log into the Network Manager; and once the call is accepted the connection is established. During the duration of the call all messages that are passed between the remote subscriber and CIP use the user data portion of the X.25 packet.

The following set of functions are available to the CIP subscriber:

- Examine and modify port parameters (trunk & standard ports)
- Examine and modify alarm routing parameters
- Examine and modify node, LP, and port statistics parameters
- Examine statistics
- Enable or disable a port
- Set an alternate N7900 ID
- Enable and disable load generation
- Enable and disable packet delay measurement
- Enable and disable packet trace
- Request operational node list
- Request initiation of database recovery
- Request initiation of file transfer between N7900s
- Request to reboot a node
- Examine the event log
- Request to re-IPL an N7900
- Open a database file
- Add a record to the database file
- Delete a record from the database file
- Request selective load of LP
- Enable and disable port loopback
- Create and remove PVCs
- Enable and disable various test services

Unlike the network operator at the console of the N7900 who can use the pop up and pull down menus and sundry screens, there is no direct operator interface when using CIP. Once the call is accepted, the subscriber would request some operation or function to be performed by embedding the function in the user data packet and sending the packet to the CIP process. CIP would break down the packet, identify the function and parameters, and then return either an error message or the resulting data as appropriate from CIP. Additional message handling processes come into play via the Message Services Process (MSP). MSP is the mechanism from which and to which all messages from CIP are sent and received by the X.25 subscriber.

CIP is a single threaded task and, since this is the case, only one connection is permitted to this process. If CIP has been requested to perform some operation and it has not completed that function and you try to perform another request an error message will be returned.

In addition to communicating with the remote subscriber and with other processes, CIP can also communicate with the PPM processes in other nodes. Some request for data can only be supplied via the local node. Another task in the local switching node referred to as the Network Process Manager (NPM) handles all request from CIP or the N7900. The NPM also performs the following functions:

- Receives and processes the messages from the N7900 Network Manager
- Sends messages to the locally attached N7900 Network Manager
- Sends messages to the NPM process on a node with a locally attached N7900 Network Manager
- Performs the "re-IPL node" function
- Monitors the locally attached N7900 Network Manager status via the PIA interface

SECTION 6

HOT STANDBY SWITCH

Network Redundancy

Total network redundancy, uptime, and availability are provided with AMNET's Hot Standby Switch. The Hot Standby Switch is designed for the modern data communications environment to facilitate the task of network management and recovery. This rack mount switch provides you with the ability to switch, monitor, test, and automatically restore operations to your network within 90 seconds.

Completely modular, this switch can grow with your network needs in capacity, function, and performance. A variety of channel interface cards are available, for example, V.35, RS-530, RS-232-C, and RJ-45. Organized in racks, the Hot Standby Switch design allows any combination of channel interface cards within the same rack. Each rack in the Standby Switch contains 16-slots for channel interface cards, one slot for a power supply and one slot for a control card. For critical applications a redundant power supply is optionally available with the Hot Standby Switch.

The Hot Standby Switch provides automatic redundancy and recovery operation of the Network Switching Nodes in the event of a nodal failure. Two nodes, with identical configurations, are connected to the network through this automatic A/B switching device. One node is active, the other node is in Hot Standby or ready state. The nodes are also connected to each other outside the A/B switch, exchanging status messages. This exchange of messages is another AMNET supplied driver. If the active node should fail, the standby node with all of the X.25 links, trunks, and access ports will now appear to the network as the active node and will automatically bring itself into the network. Figure 6-1 shows a typical configuration that uses the Hot Standby Switch.

In-process subscriber calls at the node will have to be re-established, but in-transit traffic will automatically be re-routed through other nodes. If the standby node should fail while the active node is on-line, an alarm (event) will be generated and forwarded to the N7900.

Hot Standby Switch Configuration

The Hot Standby Switch is composed of one primary rack and up to 63 additional expansion racks providing a total of 1,024 A/B switching channels. The Standby Switch can operate in two (2) modes: manual operations and automatic operations. In the manual operating mode, the operator performs all channel switching functions by using switches on the front panel of the rack.

In the automatic operations mode the default mode of operation for the N7400 or N7500 access to the Master Control Card is activated by an RS-232-C port. This port is located on either the N7500 (primary node) or N7500 (secondary node) and is connected to the Master Control Card. The Master Control Card is located in the primary Standby Switch and is interconnected with the Extension Control Cards located in each A/B extension rack.

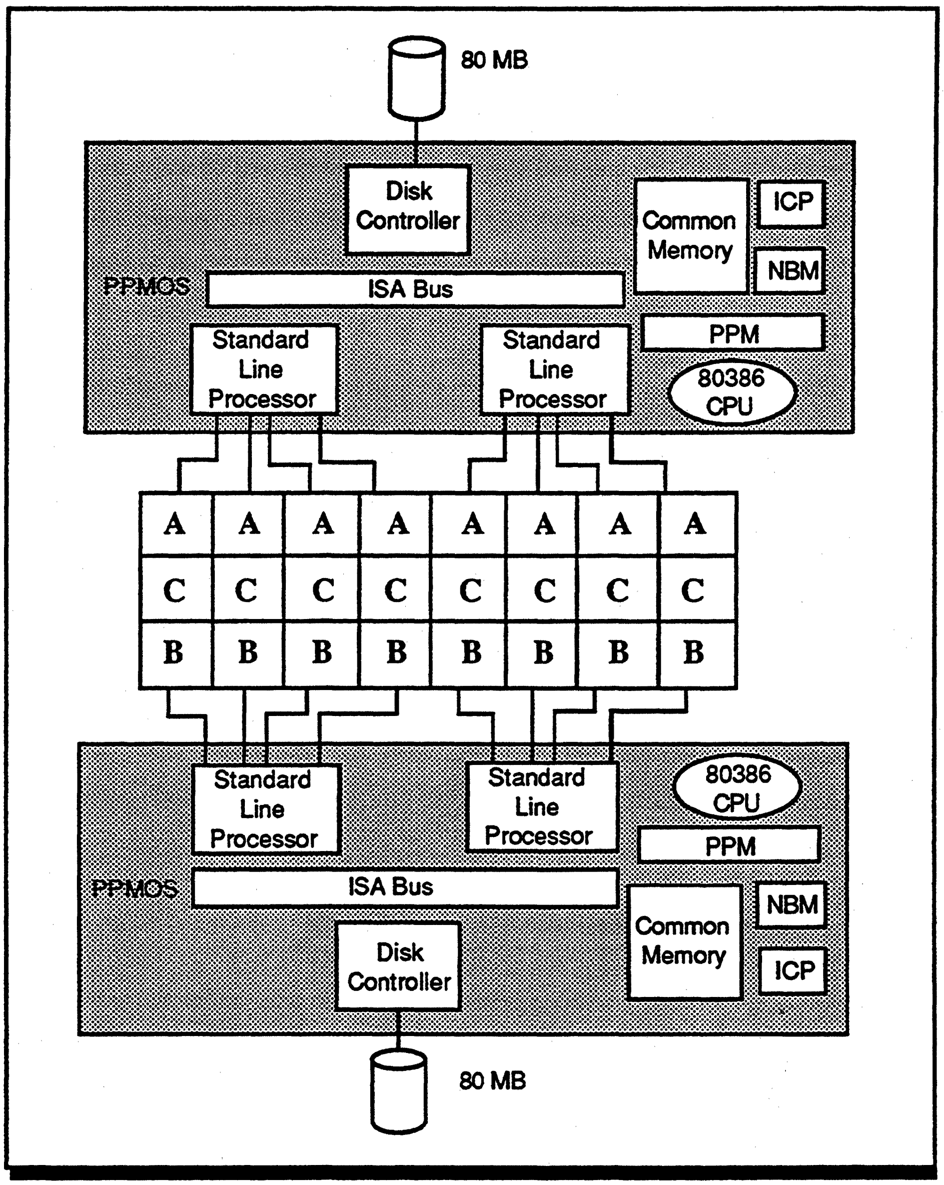

Figure 6-1 Hot Standby Switch

This provides for the automatic switching of all ports connected to AMNET's active node with all ports on the Hot Standby node. AMNET's N7900 can also be included in this operation. The only modification is the introduction of one (1) additional PIA card and cable into the N7500. Figure 6-2 illustrates the capability of a shared N7900 Network Manager between the active node and the Hot Standby node.

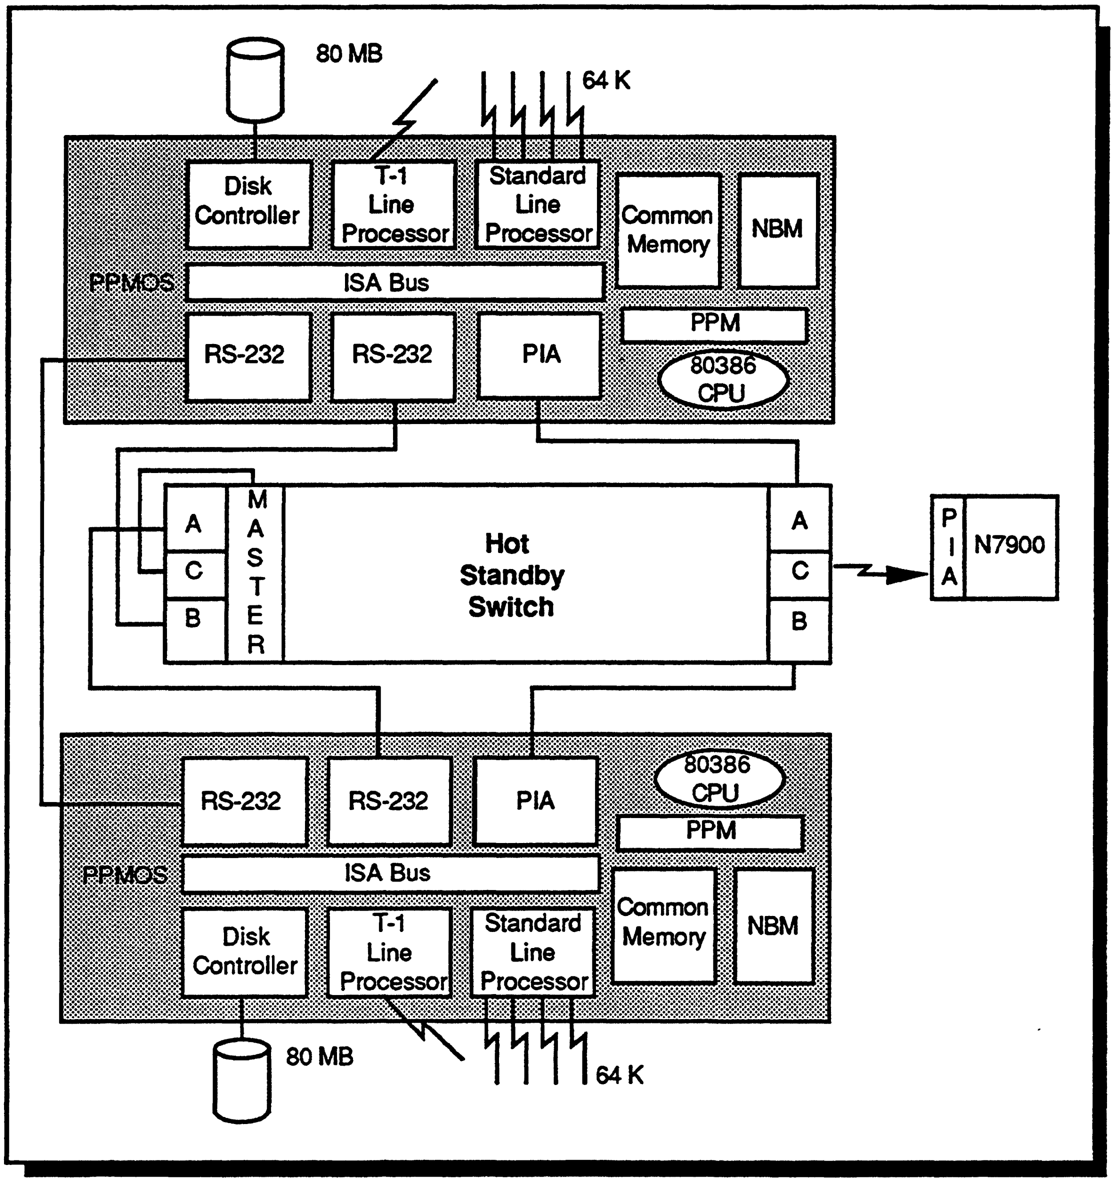

Figure 6-2 Switched N7900 Network Manager

SECTION 7

PRODUCT SPECIFICATIONS

Product Specifications

AMNET's Series of network switches includes our powerful Nucleus 7400 and the Nucleus 7500. The NSNs capabilities and features are highlighted by their flexibility, performance, and simplicity as follows:

- Dynamic least delay routing directs traffic over the best available path for the highest level of network performance
- Network trunk speeds up to 2.048 Mbps with a sustained system throughput of over 1,200 packets per second for quick user response
- Automatic adaptive routing around failures and congestion
- Packet sizes up to 4,096 bytes allow efficient use of host resources
- Integrated multiprotocol PAD support for the most popular data processing environments, eliminating the need to alter your installed hardware base
- Multiple trunk protocols supported for optimum use of internodal resources
- Automatic load balancing for efficient use of network resources
- User friendly, menu-driven configuration and monitoring capabilities for quick and easy setup and operations
- Modular design to allow you to tailor your networking requirements, growth, and performance characteristics
- State-of-the-art, industry standard hardware platform operation for ease of maintenance
- Remote configuration and download for ease of management and control

General Characteristics

- Up to four (4) software configurable ports per LP card
- Maximum aggregate throughput of 256,000 bps per Standard LP card
- Maximum packet size supports up to 4,096 bytes
- Sustained throughput rate of 1,200 packet per second (pps)
- Switched Virtual Circuits and Permanent Virtual Circuits supported
    - Supports up to 264 simultaneous VCs on N7400
    - Supports up to 2,000 simultaneous VCs on N7500
- CCITT X.121 addressing supported
- 25 - 40 calls per second (Configuration dependent)
- RS-232-C (CCITT V.24/V.28) interface
- Optional line interfaces available e.g. RS-530 and V.35
- Operating configurations loaded upon power-up and stored on disk
- On-line 'Help' commands/functions

X.25 Interface

- Compatible with 1984 CCITT Recommendations for X.25
- Configurable as DTE or DCE
- Line speeds up to 2 Mbps on each synchronous or trunk port
- X.25 facilities supported:
    - Closed User Group
    - Fast Select
    - Reverse Charging
    - Default Calling Address
    - Local Charging Prevention
    - Window and Packet Size Negotiation
    - Flow Control and Throughput Class Negotiation
    - RPOA Transit Network Selection
    - Network User Identification
- Extensive Call Control capabilities:
    - Hunt Group
    - Group Send Facility
    - Call Redirection
    - Priority Request
    - Calls Barred (Incoming and Outgoing Calls)

Asynchronous Interface

- Baud rates from 50 to 38,400 bps
- Local login password for PAD users
- Autobaud (1,200 to 19,200 baud) and autoparity supported
- Dedicated line, dial-in, and dial-out connections
- RTS/CTS hardware flow control
- CCITT X.3 (Xox/Xoff...) parameters implemented
- CCITT X.28 and X.29 implemented
- Downloadable Profiles
- Welcome banner on a per port basis
- Abbreviated addressing

- Auto disconnect with configurable timer
    - If no login
    - If no command after login
    - If no data after command
    - If clear received
    - If clear sent

SNA/SDLC Interface

- Up to 64 Kbps per port for frames up to 280 bytes
- SNA/SDLC ports can operate as DTE or DCE
- Data link control: SDLC full-duplex mode at Level I and half-duplex at Level II
- Configurable SDLC frame size
- Maximum frame size of 4,105 bytes
- Configurable station addresses
- Point-to-point operation for TPAD connections
- NRZ or NRZI supported
- Autocall on both HPAD and TPAD
- QLLC/NPSI compatible for HPAD and TPAD connections
- Four modes supported:
    - HPAD
    - TPAD
    - NPAD (Host-to-Host)
    - XPAD (Transparent)

Standard Line Processor Characteristics

- Powerful Intel 80186 Processor
- 512 Kbytes of RAM memory
- Four (4) software selectable ports per card (any combination of asynchronous or synchronous)
    - Supports SDLC/SNA, trunk, X.25, and asynchronous protocols
    - Provides direct memory access
    - Ports can be cabled for DTE or DCE
    - Provides NRZI or NRZ data encoding

- 256 Kbps total capacity (e.g., 4-ports @ 64,000 bps, 2-ports @ 128,000 bps, or 1-port @ 256,000 bps)
- Dedicated 16 Kbytes of shared memory per card
- A 16 Kbyte EPROM for initialization and control operations
- Provisions for an additional 4-port asynchronous Daughter card

Standard Line Processor Card Options

- 4-port Asynchronous Daughter card
    - Performance obtained from the attached LP card
    - Transparent bus operations
    - Available on any Standard LP card
- Line Interface available
    - RS-232-C
    - RS-530
    - V.24
    - V.35

T-1 Line Processor Characteristics

- On-board T-1 CSU and DSU requires no external equipment
- T-1 available with DSX-1 interface for microwave, satellite, or external CSU
- Fractional T-1 services supported
- Four (4) DS0 channels per card with DMA access on input and output
- Powerful Intel 80186 Processor
- 512 Kbytes of RAM memory
- Dedicated 16 Kbytes of shared memory per card
- Provisions for a 4-Channel Expander Card and 2-Port Sync/Serial Card
- Overall characteristics of the T-1 LP card is as follows:

- Line Rate	1.544 Mbps
- Line Code	AMI or B8ZS
- Framing	D4 or ESF
- Indicator	Rx Frame Sync Light Emitting Diode
- Clocking	Master or Slave
- Connector	Standard RJ-48C jack connector
- Registration	FCC Part 68

E-1 Line Processor Characteristics

- CCITT G.703/G.704 Recommendations compliance
- Up to 30 channels E-1 (CEPT) standards for European Service supported
- Four (4) DSO channels per card with DMA access on input and output
- High-performance Intel 80186 Processor
- 512 Kbytes of RAM memory
- Dedicated 16 Kbytes of shared memory per card
- 4-Channel Expander Card and 2-Port Sync/Serial Card provisions
- Overall characteristics of the E-1 LP card is as follows:

- Line Rate	2.048 Mbps
- Line Code	HDB3
- Framing	G.704 Recommendation
- Indicator	Rx Frame Sync Light Emitting Diode
- Clocking	Master or Slave
- Connector	BNCs for 75 ohm Receiver and Transmitter coaxial cables

T-1/E-1 Line Processor Card Options

- 4-Channel Expander Card
    - Provides four (4) additional logical channels per T-1 or E-1 LP card
- 2-Port Sync/Serial Card
    - Two (2) independent high-speed ports that are software configurable
    - Programmable speeds of Nx56/64s up to 1.536 Mbps for T-1 facilities or 1.920 Mbps (E-1) supported
    - Data switching of incoming T-1/FT-1 or E-1 facilities provided
    - DCE or DTE mode operation
- Line Interface available on 2-Port Sync/Serial Card
    - RS-422
    - V.35

APPENDIX A

CALL VALIDATION PROCESS

Call Verification

The call verification process is a relatively simple operation. This operation is performed either by the N7900 Network Manager or an N7500 with ICP installed. The call verification process verifies several items. When the X.25 Level III process receives an SVC call set-up request, it verifies the following items:

- Length of call request packet - greater than five (5) octets
- A valid General Format Identifier
- Valid length of calling subscriber identification
- Valid length of called subscriber identification
- Length of facility field less than or equal to 63 octets
- For Group Send operations the group send facility code must be in the facility field
- Valid calling subscriber identification verified against the database

If these items are valid, then the SVC call is ready to be established. To better understand the entire network flow, we will illustrate this process by having an Async PAD subscriber place a SVC call and follow its flow to its logical conclusion (call establishment).

Async SVC Call Establishment

The subscriber would enter his/her subscriber ID usually followed by a password. The password security is established by the network administrator. In some cases this is optional, but it can be established that the password is mandatory. This would be followed by entering the called address of the destination DTE. Here the subscriber could use the abbreviated address or mnemonic address. Therefore, the call setup packet would contain the following information:

- Source node
- Destination node
- Logical channel number
- Calling subscriber ID
- Called subscriber ID or Group ID

This packet information is at the source node and is moved into the X.25 Level III process which verifies the packet as being a legal packet and where it should be sent. The Level III provides us with the switching and routing information. This packet can be forwarded either to an N7900 Network Manager or to an N7500 running the ICP process. This is established by the network administrator as part of the node record. The network administrator can also select a primary and secondary destination address for call processing and configuration download operations.

Once the packet reaches the N7900 the aforementioned call verification process is executed. Again, the calling/called address is verified as well as the subscriber IDs, and Group IDs with the information resident in the database. If it passes we send back a call acknowledgement to the Level III process with the LCN that will be used once this call is established. If it is rejected we send back a negative call response message with the reason for rejection. (The call reject codes are outlined in the next sub-section.)

Now we can send a packet to the destination node. This is accomplished via the *Network Process* which provides the communications capabilities among the PPM processes on different nodes. Again we will use the Level III process, but this time it is the destination node which initially receives the packet. We check that the Level II is up and functional on the destination node. If Level II is up, we reformat the packet to an incoming call request packet and wait for a call accepted confirmation from the local DTE.

When the call accepted packet is received at the destination node, we establish a single connection via the Level III process and send the call accepted packet back to the source node. This is accomplished through the *Network Process*. At the Level III process at the source node we establish a single connection, reformat the call accepted packet into a call connected packet and enter the data transfer state. This is the entire call establishment process between a source and destination node.

Call Reject Codes

There are only three (3) call reject reason codes generated by the N7900 or N7500. But within each of these codes are multiple diagnostic codes. The following is a list of the cause code/diagnostic code combinations:

03 - Invalid Facility

- 65 - Facility not allowed

11 - Access Barred

- 132 - No incoming access
- 133 - No outgoing access

19 - Local Procedure Error

- 67 - Invalid called address
- 68 - Invalid calling address
- 128 - File access failure
- 129 - No subscribers in group
- 130 - Invalid logical channel number
- 131 - Invalid closed user group

APPENDIX B

HIGH-SPEED OPTIONS

High-Speed Line Processor Cards

High-speed T-1 and E-1 Line Processor cards are now supported on the N7400 and N7500 Network Switches. Fractional T-1/E-1 (FT-1/FE-1) speeds for internodal trunk traffic is also supported. The T-1 LP cards operates at 1.544 M bits/sec and the E-1 LP at 2.048 Mbps. Like the Standard LP cards, these cards, support a shared memory area where the workstation and the LP can access this common memory area.

Two (2) optional cards are available with the T-1/E-1 LP: a 2-port Sync/Serial card and a 4-Channel Expander card. The 2-port Sync/Serial card allows you to perform T-1 multiplexing services, that is, it provides for the transparent carrying of synchronous data from other devices over the T-1/E-1 line. The 4-Channel Expander card, used in conjunction with the T-1/E-1 card, provides four (4) additional DS0 channels.

There are three (3) different LP models available. The N7450-001 is a T-1 LP card that has an integrated (on-board) CSU. This provides you with the ability to directly connect to a T-1 or FT-1 line. The second model is the DSX-1 model (N7450-000) which does not have an on-board CSU. It is intended for connection to an external CSU connected to a T-1 line or for connecting to a local T-1 device with a DSX-1 interface like another T-1/E-1 LP card. And the last card is a E-1 (CEPT) LP card (N7450-002) for connection to an E-1 CEPT line, or for connecting to a downstream E-1 device.

Fractional T-1/E-1 speeds are supported by using the Alternate Mark Inversion (AMI) or Bipolar Eight Zero Code Substitution (B8ZS) line coding techniques. For example, a group of DS0 can be used for either Nx56 Kbps or Nx64 Kbps operations, where N is a value between 1 and 30. The AMI line coding technique employes a ones density method to carry clock signals which results in the utilization of a 56 Kbps bandwidth over each 64 Kbps DS0. Sometimes this is also referred to as the AMI with Zero Code Substitution (ZCS). The second technique employed is Bipolar Eight Zero Code Substitution (B8ZS). This method does not use the ones density method for clocks which results in the full use of the 64 Kbps of available bandwidth for each DS0.

Just like the Standard LP cards, each T-1/E-1 LP has a shared memory area. The addresses for the shared memory area is constant for either type of LP card with the noted exception that the I/O base addresses are different since we could not establish the identical I/O base addressing scheme used by our Standard LPs. The I/O base addresses provides us with a means of determining what type of LP card is physically located in the node. Table B-1 illustrates the Standard LP card I/O base address and its associated shared memory region whereas Table B-2 shows the new T-1/E-1 LP cards base address. (All values in both tables are in hex format.)

LP Number	Shared Memory Base Address	I/O Base Address
0	C000:0	240
1	C400:0	250
2	C800:0	260
3	CC00:0	270
4	D000:0	280
5	D400:0	290
6	D800:0	2A0
7	DC00:0	2B0

Table B-1 Standard LP Address Assignment

LP Number	Shared Memory Base Address	I/O Base Address
0	C000:0	0100
1	C400:0	0900
2	C800:0	1100
3	CC00:0	1900
4	D000:0	2100
5	D400:0	2900
6	D800:0	3100
7	DC00:0	3900

Table B-2 T-1/E-1 LP Address Assignment

As you can see the T-1/E-1 LP will have shared memory locations in the same address space as the Standard LPs, but its I/O address space is different. The software will determine what type of LPs are configured on the node via the I/O base addresses and configure its shared memory accordingly.

Through the use of the 4-Channel Expander (N7450-020) card you can increase the number of DS0 groups in use on the T-1/E-1 line. This card does not connect directly into the T-1/E-1 line but is connected via an over-the-top connector onto the T-1 or E-1 LP card.

The T-1/E-1 LP provides the Extended Superframe Format (ESF) Facility Data Link (FDL) support in accordance with AT&T Technical Publication 54016, Requirements for Interfacing Digital Terminal Equipment To Services Employing The Extended Superframe Format, May 1986. This publication describes a procedure to collect real-time T-1 circuit performance and alarm information. Of the total 8 Kbps available for framing in a T-1 line, the ESF uses 2 Kbps for framing, another 2 Kbps for data error checking (CRC6), and the remaining 4 Kbps for the Facility Data Link which is used as a maintenance communications channel.

The T-1/E-1 LP will monitor the following errors: first, ESF error event (i.e., the receipt of a bad CRC6); second, Errored Seconds (ES); and third, Unavailable Seconds (UAS). An Errored Seconds is defined as a second with one or more ESF error events that is multiple data error checking. When the rate of ESF error events is very high, in one second, than the line is considered unavailable is logged in Unavailable Seconds. This ES and UAS data is available for the T-1 carrier to retrieve through the FDL channel for monitoring and maintenance of the T-1 line.

The aforementioned T-1/E-1 LP cards are configured as primary LPs. As a primary mode LP the card is programmable to obtain timing from the on-board timers or it derives the timing source from the network. Each T-1/E-1 LP card supports four (4) DS0s. All the existing LP and port statistics are supported for the T-1/E-1 LPs as well as unsolicited alerts. Alert history is also maintained by the T-1/E-1 LP card. Additionally, ESF statistics are available for the CSU and the DSX-1 LP. The ESF statistics available are as follows:

- One (1) hour TELCO data
- Twenty-four (24) hour TELCO data
- One (1) hour User data
- Twenty-four (24) hour User data
- Reset User registers

A summary of some of the key features of the T-1/E-1 LP card are:

- T-1/E-1 LP is supported in both the N7400 and the N7500
- The T-1/E-1 LP card can coexist with the Standard LP card in the same node
- Eight (8) LP of any mix are supported
- Each T-1/E-1 LP supports four (4) ports in any combination
- Standard Trunk (LAPB) protocol and Fast Trunk protocol support is available
- For multiple port operations the aggregate total line speed is 1.544 Mbps for T-1 line and 2.048 Mbps for E-1 line
- Each port is capable of supporting fractional T-1/E-1 speeds that are multiple of a DS0
- Up to 24 groups of DS0 are supported for T-1 operations and up to 30 groups for E-1
- Up to 16 Sync/Serial cards are supported per T-1/E-1 LP card
- Up to seven (7) 4-Channel Expander cards are supported per T-1/E-1 LP

T-1 Line Processor Characteristics

• Line Rate	1.544 Mbps
• Line Code	AMI or B8ZS
• Framing	D4 or ESF
• Indicator	Rx Frame Sync Light Emitting Diode
• Clocking	Master (internal) or Slave (external)
• Connector	Standard RJ-48C jack connector for CSU model or RJ-45 for DSX-1
• Registration	FCC Part 68
• Optional	4-Channel Expander card and 2-Port Sync/Serial card

E-1 Line Processor Characteristics

• Compliance	CCITT G.703/G.704 Recommendations
• Line Rate	2.048 Mbps
• Line Code	HDB3
• Framing	G.704 Recommendation
• Indicator	Rx Frame Sync Light Emitting Diode
• Clocking	Master or Slave
• Connector	BNCs for 75 ohm Receiver and Transmitter coaxial cables
• Optional	4-Channel Expander card and 2-Port Sync/Serial card

T-1/E-1 Line Processor Card Options

- 4-Channel Expander card
  - Provides four (4) additional logical channels per T-1 or E-1 LP card
- 2-Port Sync/Serial card
  - Two (2) independent high-speed ports that are software configurable
  - Supports programmable speeds of Nx56/64s up to 1.536 Mbps for T-1 facilities or 1.920 Mbps (E-1)
  - Provides data switching of incoming T-1/FT-1 or E-1 facilities
  - Operates in either DCE or DTE modes
- Line interface available on 2-Port Sync/Serial card
  - RS-422
  - V.35

T-1/E-1 Channel Mapping

All DS0 channels are controlled by the T-1/E-1 LP card. The establishment of the various channel or connection mapping assignments are performed by the network administrator. This channel mapping process can be thought of as placing into service all the DS0 channels on the T-1 or E-1 LP card. Twenty-four (24) DS0 channels are available on the T-1 LP card and up to thirty (30) DS0 channels on an E-1 line.

The channel mapping process consists of initially configuring the four (4) DS0 channels on the Primary T-1 or E-1 LP card. Then, mapping any additionally 2-port Sync/Serial cards followed by any attached 4-Channel Expansion cards.

Sync/Serial Card

The 2-Port Sync/Serial card is supported via the T-1/E-1 LP card. The Sync/Serial card is connected to the T-1/E-1 LP card with an over-the-top ribbon cable. Sixteen (16) Sync/Serial cards can be connected to a T-1/E-1 LP or up to 96 Sync/Serial cards per node. However, the full number of Sync/Serial cards supported depends on the number of available slots in the workstation. Each Sync/Serial card supports two (2) independent synchronous ports for connecting a DTE to a T-1/FT-1 line. Both ports will operate at speeds that are multiples of a DS0s and will support either a RS-422 or V.35 line interface.

The Sync/Serial card does not use the shared memory scheme but standard control signals are supported as: Request to Send, Clear to Send, Receiver Ready, Data Set Ready, and Data Terminal Ready.

Sync/Serial Card I/O Address Assignment

The 2-Port Sync/Serial card does not have any shared memory requirements but it does need I/O addresses. There are 16 I/O base addresses assigned for each LP numbered 0 - 3 . This provides the first four (4) LPs to support up to a 16 Sync/Serial cards each, for a total of 64 Sync/Serial cards. For LPs numbered 4 - 7, eight (8) I/O base addresses are assigned to yield an additional 32 Sync/Serial cards, for a total of 96 Sync/Serial cards per node.

Support of the 2-Port Sync/Serial cards in a node depends on the number of available slots and the I/O decode capability of any other controllers installed on the node. Table B-3 provides the LP number and the Sync/Serial card I/O Base Address assignment. (All I/O base address values are in hex format.)

If another controller decodes only ten (10) I/O address bits then it will not be possible for it to coexist with the Sync/Serial card. For example, the Archive VP402 Tape Drive interface card may not coexist with the Sync/Serial cards assigned for LP 0 and LP 3, a Game I/O adapter card with the Sync/Serial card for LP 2, LPT1 parallel printer card with the Sync/Serial cards for LP 6 and LP 7, and LPT2 parallel printer card with the Sync/Serial card for LP 5.

Sync/Serial Card Number	LP 0	LP 1	LP 2	LP 3	LP 4	LP 5	LP 6	LP 7
0	0140	0180	0200	0300	0500	0640	0340	2340
1	0540	0580	0600	0700	0D00	0E40	0740	2740
2	0940	0980	0A00	0B00	1500	1640	0B40	2B40
3	0D40	0D80	0E00	1F00	1D00	1E40	0F40	2F40
4	1140	1180	1200	1300	2500	2640	1340	3340
5	1540	1580	1600	1700	2D00	2E40	1740	3740
6	1940	1980	1A00	1B00	3500	3540	1B40	3B40
7	1D40	1D80	1E00	1F00	3D00	3E40	1F40	3F40
8	2140	2180	2200	2300				
9	2540	2580	2600	2700				
10	2940	2980	2A00	2B00				
11	2D40	2D80	2E00	2F00				
12	3140	3180	3200	3300				
13	3540	3580	3600	3700				
14	3940	3980	3A00	3B00				
15	3D40	3D80	3E00	3F00				

Table B-3 Sync/Serial Card Address Assignment

Node Initialization and Operation

During the IPL of a node, *lnode* is responsible for the reset of the LP hardware, the download of the LP software, and the setup of the LP shared memory for the PPM/LP communications. *lnode* supports the Standard LP cards, the T-1/E-1 LP cards, and the 2-port Sync/Serial cards. The first operation performed by *lnode* is to build a table of installed LPs on the node. This is accomplished by disabling the shared memory region for any T-1/E-1 LP card. Next the workstation memory is scanned to determine where the Standard LPs are located by looking at eight (8) pre-defined shared memory areas. Thereafter, the shared memory for all possible remaining LPs are enabled using the T-1/E-1 LP I/O base addresses.

Once this table is constructed, we know the type of LP and where it is located in the node. We issue a reset and then download all installed Standard LPs with our Standard LP software from the floppy or hard disk (i.e., FEP.EXE). Next we reset and download all installed T-1/E-1 LPs first with a bootstrap code and then with the T-1/E-1 LP software (T1FEP.EXE). Next the PPM software can now arrange the shared memory of the existing LPs in a linked list for standard operations.

The final procedure is to build a table of any installed Sync/Serial cards by examining all possible 96 Sync/Serial card I/O addresses. All Sync/Serial cards that are detected are reset and loaded with their Sync/Serial program (i.e., LCA.BIN). This completes the IPL procedure and the PPM software is now loaded into the workstation and executed. Once the trunk lines are established the node is recognized by its adjacent nodes and it is broadcasted to the rest of the network.

APPENDIX C

INTEROPERABILITY

Interoperability

AMNET's products span the gamut from our Nucleus 7200 (N7200) Multiprotocol Switching PAD to our powerful N7400 and N7500 Network Switches culminating with our N7900 Network Manager. The N7200 performs operations as an X.25 switch, X.25 gateway, and multiprotocol PAD. The Nucleus 7200 can be located anywhere in the network and can be controlled remotely. The N7200 LP cards support up to four (4) synchronous or asynchronous ports expandable in 4-port increments. Like the NSNs, each port on the N7200 is software selectable on a port-by-port basis and supports the following protocols:

- Asynchronous (X.3, X.28, X.29)
- SNA/SDLC
- 3270 BSC
- 2780/3780 BSC RJE environments

The N7200 also supports our 4-port Async only Daughter card. When the combination of our 4-port Daughter card and the LP cards are utilized, a total of 64-ports are available. Thirty two (32) of these ports are Async ports and the rest can be selected as access or trunk ports depending on your requirements. The N7900 Network Manager supports all the NSNs in the network as well as any N7200 - total network support for any node in the network.

Operators can access and control the N7200 directly from the N7900 Main Menu selection screen as illustrated in Figure 3-3. By selecting *N7200 Remote Control* the operator can access an N7200 anywhere in the network. The operator must answer some questions to acquire control of the N7200. For example, the operator must have a remote control subscriber ID number and password to place the call to the remote N7200. Also, the operator must know the remote N7200 password to gain access into the N7200. Several levels of password protection are provided to insure subscriber and network security.

Remote operations can be controlled directly from the display console of the N7900 or remote control can be accomplished via a PAD connected with an ANSI terminal. This ANSI terminal defaults to a VT100 class device and, therefore, it must be on an Async PAD port. Once the operator has provided the appropriate responses the menu will be displayed as illustrated in Figure C-1.

If the remote N7200 has a display monitor attached, the status line of the display will show the following message:

****Under Remote Control. Keyboard blocked by ID = 9999990018****

The operator now has total control of the N7200. The main menu displayed is exactly what the local operator would see when the menu selection is initiated. Changes can be made permanently and written to the remote N7200

```
AMNET N7200 Version 2.xx.yy

MAIN MENU

A   Set Concentrator Parameters
B   Set Port Parameters
C   Subscriber Management
D   Status Display
E   On-line Control
F   Remote Control
G   Change Password
Select:
```

Figure C-1 N7200 Main Menu

diskette or the modifications can be temporary and kept in the N7200 RAM memory. If the changes are temporary and the N7200 is re-IPLed then the original settings would be initiated from the local diskette.

This operation is part of the SVC test services. The called address displayed on the remote N7200 monitor reflects the fact that it is part of the test services. For example, with the ID = 9999990018 the first six (6) nines (i.e., 999999xxxy) signify that the ID is part of the test services function. The next three (3) digits (xxx or 001) are the node ID number that is node one (1), and the last digit eight (8) signifies that the SVC call requested use of the N7900 *consol* process.

Total Network Interoperability

Figure C-2 shows a typical network incorporating AMNET products and some of the applications they support. AMNET's N7900 Network Manager provides total support and control for PADs and network switches.

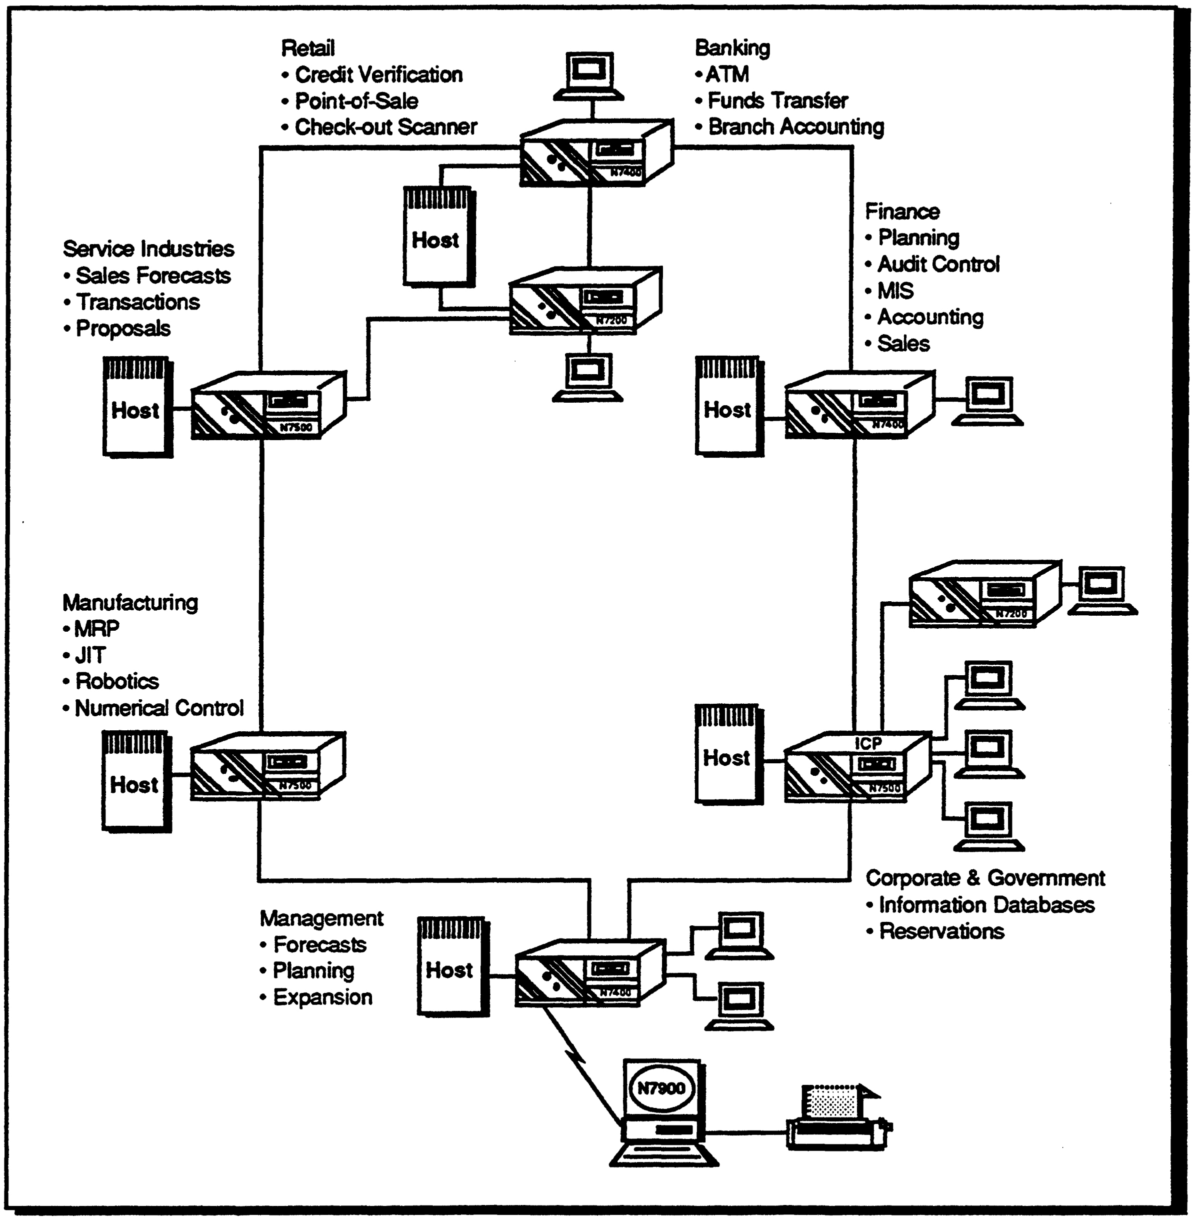

Figure C-2 AMNET's Backbone Network

APPENDIX D

IPL OPERATION

IPL Routine

The Standard LP card is composed of a dedicated 80186 processor with 512 Kbytes of RAM memory, dual ported shared memory, and ROM memory. The ROM memory contains the code for initialization; i.e., to functionally self-test the LP card during start up operations; load operations used to downline load the appropriate LP operating system (FEPOS), PAD software/applications, and X.25 trunk protocols, tasks, and lastly control primitives. The cold start routine is entered by a hardware reset, initiated by a power up condition or by a reset command on the workstation.

The 16 to 128 Kbytes of shared memory (dual-ported memory) is the primary interface between the workstation and the LP card. The status and command bytes in this area are constantly updated. Some of the status bytes determine the number and type of error code received, line status registers, and free buffer pool status.

The cold start routine executed by the LP ROM performs a self-diagnostic test. The self-diagnostic completes the following tests:

- First test the LP processor internal registers, flags and operational codes.
- Next, all memory is cleared by performing a series of reads and writes to memory. Once the memory is initialized, a memory diagnostic is executed to further test the memory.
- All timers are initialized and tested.
- Initialize and test the DMA controllers.
- Perform all primitive operations with regard to controlling the shared memory region.

If any one of these tests fails an error code is returned. When the workstation tries any further operation to the LP it is informed about the error condition. On successful completion of the above tests, the control of the LP is transferred to an idle condition and waits for a command from the workstation.

The four (4) communications ports of the LP card are brought out through the rear panel of the workstation platform by a 62-pin connector mounted at the rear of the LP card. A special 4-port splitter cable comes standard with each LP card and provides four (4) standard (RS-232-C) DB-25 connectors. The splitter cable is available in various configurations to support a myriad of DTE or DCE port interfaces.

The IPL task (ipltsk) performs the function of loading and establishing the PPMs and the LPs on the node. The operation of this task is actually controlled by the PPM itself with the N7900 software operating in a slave capacity.

A description of the operations of the IPL process follows:

1) PPM requests IPL with EDBCC message.

2) *ipltsk* acknowledges start of IPL with ADBCR message.

3) PPM requests configuration table with XIPLC message.

4) *ipltsk* sends back node configuration list (NCNFL).

5) PPM takes intersection of node configuration list and the actual processors that are there and requests FEP load images (XCDT) and PPM load images (XMDT).

6) *ipltsk* sends load image end (EXLIF).

7) PPM sends end of downline load phase command (DLLC).

8) PPM sends request for LP parameters (IPL2).

9) *ipltsk* sends port type count record (CNTREC).

10) *ipltsk* sends alternate node message (SETANO0).

11) *ipltsk* sends network defaults (SYSREC).

12) *ipltsk* sends port records to PPM (TRKREC for trunks, and OTHREC for data ports.)

13) *ipltsk* sends port records to LPs (IPLPAR).

14) *ipltsk* sends end of IPL Phase 2 message (IPLCMP).

15) PPM sends request for Phase 3 message (IPL3).

16) *ipltsk* sends alarms routing messages (ALMMOD).

17) *ipltsk* sends node statistics parameters (MBSMOD).

18) *ipltsk* sends end of node statistics parameters (MBSEND).

19) *ipltsk* sends LP/port statistics parameters (FPSMOD).

20) *ipltsk* starts PVC generator task and informs it to generate all PVCs that have both ends on this node.

21) *pvcgen* task sends PVC creation messages (PVCREC).

22) PPM acknowledges PVC creation request with PVCCAK (PVC was created) and a negative response with a PVCCNK message (PVC was not created).

Figure D-1 illustrates a workstation and LP combination. A detailed description of the operations between the workstation and LP follows.

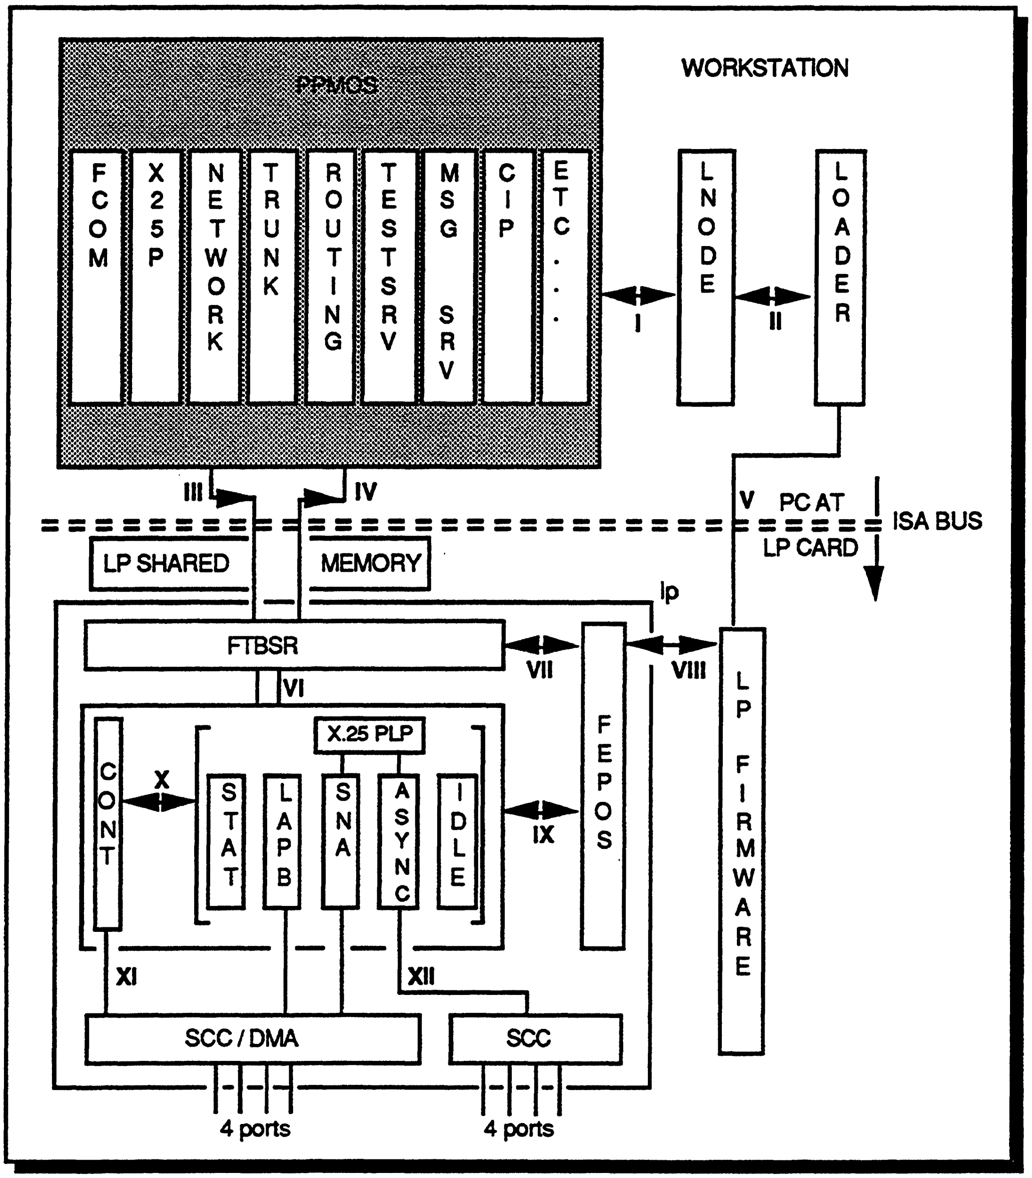

Figure D-1 Standard LP Port Interface

I. *Interface between PPM and LNODE (IPL only).* The LNODE.EXE module runs under DOS.

II. *Interface between LNODE and LOADER (IPL only).* The LNODE module uses the LOADER to load the LP software as a single module on the LP card.

III. *Interface between PPM and LP (PPM -> LP).* The interface is done through a transmit queue located in the shared memory region. The access to this queue is controlled and the buffers are obtained from the shared memory free buffer pool. The *IDLE* task in the LP watches this queue and invokes *FTSBSR* when the queue is not empty (i.e., there is an incoming message waiting for the LP). No hardware interrupts are used. The PPM writes into the queue through the *FCOM* task when:

- There is a buffer available in the shared memory, and
- The port is write enable

IV. *Interface between LP and PPM (LP -> PPM).* The interface is done through a receive queue located in the shared memory. The access to this queue is also controlled. The *FCOM* task in the PPM regularly polls this queue. The *FTBSR* task stops adding to this queue when the number of buffers available in the shared memory drops to eight (8) or below for flow control purposes. When the PPM holding queue crosses a threshold, the FEP applications will go into flow control state for the line causing the problem. For example, RNR control frame is sent to the line in LAPB applications.

V. *Interface between LOADER and LP Firmware.* This interface exists during IPL only. Its purpose is to download the FEP software to the LP card.

VI. *Interface between FTBSR and all sub-tasks. FTBSR* task collects all messages from the Node Buffer Memory exchange and sends them to the PPM through the shared memory. It also extracts all the incoming messages (from the PPM) and routes them to the configured application task using a dispatch table.

VII. *Interface between FEPOS and FTBSR.* FTBSR is one of the sub-tasks running under FEPOS.

VIII. *Interface between LP Firmware and FEPOS (IPL only).* This interface exists during IPL time only. Its purpose is to start the FEPOS initialization routine when the download finishes.

IX. *Interface between FEPOS and all tasks.* FEPOS manages the sub-tasks on a priority basis.

X. *Interface between CONT (control) and other tasks.* *CONT* has two (2) functions: first, at the time of IPL, it sets the parameters of each port of the card depending on the configuration (i.e., LAPB, SNA/SDLC PAD, 3X PAD); and second, *CONT* is also the LP control routine dispatcher. Its function here is to process messages sent to the LP controller application according to their type.

XI. *Interface between CONT tasks and SCC/DMA chip.* On an N7900 request, the *CONT* sub-task checks the port signals and sends a message to the N7900 through the *FTBSR/FCOM* interface.

XII. *Interface between sub-tasks and SCC/DMA chip.* During initialization the DMA/SCC registers and interrupt addresses are set by the applications, that is, LAPB, Async, and SNA/SDLC. In normal operations, the applications use this interface to transmit and/or receive data from the ports through the DMA/SCC interrupts.

GLOSSARY

A

ACK — A communication control character transmitted by a receiver as an affirmative response to a sender. It indicates that the preceding transmission block has been received, and that the receiving unit is ready to accept the next block of the transmission.

Address — A coded representation of the destination or origin of a message.

Adaptive Routing — Routing that automatically adjusts to network changes such as altered traffic patterns or failures.

Alarms — Status messages that appear on your printer to inform you of routine network events and warn you of any abnormal situations anywhere in the network.

Algorithm — A defined set of operational steps to be taken to effect a desired calculation.

ANSI Terminal — American National Standards Institute establishes standards for telecommunications networks and devices.

ASCII — American Standard Code for Information Exchange is a seven level code set (128 possible characters) with a provision for parity that is used for data transfer in the U.S.

Async — An abbreviate form of asynchronous, a term used to describe terminals, printers, and other devices that transmit data in a discontinuous character stream. It also describes the protocol used to interconnect such devices.

Async Console — An Async terminal connected to a serial port configured as COM1.

Async Terminal — A VT100 or compatible video display terminal.

Autobaud — The default throughput class for Async PAD ports. When a PAD places a call through a port set to autobaud, the port's internal clocks are automatically set to match the speed of the PAD.

B

Banner The display area at the top of a screen that generally features product name, revision, etc.

BPAD Bisynchronous Packet Assembly/Disassembly

BSC Binary Synchronous Communications. A character-oriented protocol developed by IBM oriented towards half-duplex link operations.

Buffer A storage device or routine used to compensate for a difference in the rate of data flow, or the time of occurrence of events, when transmitting data from one device to another. A buffer permits a change of speed or interface characteristics between two (2) devices.

C

Call A request for connection between two (2) DTEs.

CCITT Consultative Committee for International Telephone and Telegraph which developed the X.25 standard for packet switching. It is a consulting committee of the International Telecommunications Union of the United Nations.

CEPT Conference of European Postal and Telecommunications administrations is a European body that makes recommendations for telecommunications practices in Europe.

CIP Computer Interface Process allows a remote subscriber to access the full complement of resources available on the N7900 Network Manager with the exception of graphics.

Cluster Controller A device that handles the remote communications processing for multiple terminals or workstations (i.e., IBM 3270-family controller).

CMM Common Memory Module is used for such transient purposes as connection records by the PPM tasks and for passing messages by PPMOS.

C

Concentrator	Any communications device that allows a shared transmission medium to accommodate more data sources than there are channels currently available to the transmission medium.
Connection	A logical path between two (2) ports over a network route that allows the exchange of data.
CPU	Central Processing Unit. A computer or the part of a computer system that interprets and executes instructions.
CSU	Channel Service Unit is an interface between the customer premises and the operating phone company. The CSU is used to terminate a T-1 circuit and provides network protection and diagnostics.

D

D-4	A T-1 framing format that complies with AT&T Technical Reference 62411 in regards to the DS-1 frame layout.
Database	The network configuration files, subscriber files, statistics and events are organized into a database for quick key access and security.
DATAPAC	A Public Data Network (PDN) supported in Canada.
D-Bit	The delivery confirmation bit in an X.25 packet that is used to indicate whether or not the DTE wishes to receive an end-to-end acknowledgement of delivery.
DCE	Data Circuit-Terminating Equipment is any device that maintains and terminates a connection between the data terminal equipment and a communications facility (i.e., modem).

D

DMA	Direct Memory Access. A method by which information can be transferred from the computer memory to a device on the bus without using the processor or CPU.
DOD	Department of Defense
DSP	Display System Protocol
DSU	Data Service Unit is a device used to interface to a digital circuit (i.e., DDS or T-1 when combined with a CSU). It performs conversion of customer's data stream to bipolar format for transmission.
DSX-1	Digital Signal Cross-Connect Level 1 refers to DS-1 signal level at the cross connect patch panel found at the Central Office.
DTE	Data Terminal Equipment is any device which transmits data to and/or receives data from a data communications system (i.e., a CPU or terminal).

E

EBCDIC	Extended Binary Coded Decimal Interchange Code is an eight-bit character code used primarily in IBM equipment. The code allows for 256 different bit patterns.
E-O-T	End of Transmission
EIA	Electronic Industries Association is a standards organization in the United States.
EIA Interface	A standardized set of signal characteristics (time duration, voltage, and current) specified by the Electronic Industries Association.
Event	Any alarm condition, error, or normal occurrence reported to the system.

F

Facility — The transmission path between two (2) or more points provided by a common carrier.

FC — Flow Control is the capability of network nodes to manage buffering schemes while handling devices operating at different rates, enabling them to talk to each other without loss of data.

FCS — Frame Check Sequence is a bit-oriented protocol, a 16-bit filed that contains transmission error checking information, usually appended to the end of the frame.

FEP — IBM's Front End Processor

FEPOS — Front End Processor Operating System is the AMNET developed operating system supporting its LP cards.

Frame — A logical transmission unit sent between data link layer entities that contains its own control information for addressing and error checking.

G

Gateway — A conceptual or logical network station that serves to inter-connect two (2) otherwise incompatible networks, network nodes, subnetworks, or devices (performs a protocol-conversion operation across numerous communications layers).

General Format Identifier — The General Format Identifier (GFI) is the first four bits of the first octet (the high-order bits) for each packet header format. The header is composed of three octets; octet two (2) is the LCN and octet three (3) is the packet type identifier.

H

HDLC	High-level Data Link Control is a CCITT specified, bit-oriented, data link control protocol on which most other bit-oriented protocols are based.
HLCN	High logical channel number
Hot Standby Node	An optional redundancy feature that consists of a duplicate node and A/B switch assembly that provide immediate switch-over in case of failure of the primary node.
HPAD	The Host-to-Terminal PAD
Hunt Group	A designated group of ports that share incoming calls according to a pre-defined algorithm.

I

ICP	Integrated Call Processing is an application program available on the N7500 which provides local call verification, call setup and overall improvement in network response time.
I/O	Input/Output is the process of transferring data from or to a computer system including communication channels, or operator interface devices and control channels.
IP	Internet Protocol used in gateways to connect networks at OSI Level III and above (normally associated with Level IV protocol TCP).
IPL	Initial Program Load is a process whereby the workstation, node, or N7900 is initialized. This usually involves the loading of operating system, task, and application programs.

I

ISDN Integrated Services Digital Network is a CCITT recommendation of operating parameters and interfaces for a network that will allow a variety of mixed digital transmission services to be accommodated.

ISO International Standards Organization which is involved in writing communications standards using the OSI Reference Model.

L

LAP The Data Link-level protocol specified in the CCITT X.25 interface standard.

LAPB/LAPD LAPB is a newer version of LAP specified for balanced applications (either node may initiate transmission, rather than one as in a master/slave arrangement); and LAPD is specified for ISDN applications and differs from LAPB in its framing sequence.

LCN Logical Channel Number which is used as an identifier for a virtual circuit between two (2) DTEs.

Link A logical circuit between two (2) users of a packet switched network allowing them to communicate; to indicate the existence of communications facilities between two (2) points.

LP Line Processor cards perform the input/output functions in the N7400 and N7500 Network Switches.

LU Logical Unit in IBM's System Network Architecture (SNA), a port through which a user gains access to the services of a network.

M

M-Bit	The More Data bit in an X.25 packet that allows the DTE/DCE to indicate there is more current packet related information to follow.
Message	A defined entity of information consisting of one or more packets.
Multiprotocol	The term used to describe a device which supports several networking protocols.
Multitasking	Generically refers to the concurrent execution of two (2) or more tasks by a computer.

N

NBM	Node Buffer Memory is a dedicated area in the NSNs to hold frame and packet information.
NCP	Network Control Program is an IBM host generated program that controls the operation of a communications controller (i.e., IBM 3725).
Network	A series of points connected by communications channels.
Node	A point where one or more functional units interconnect transmission lines; and a physical device that allows for the transmission of data within a network.
NPAD	The Host-to-Host PAD
NPSI	Network Control Program Packet Switching Interface
NRZ	Non-return to Zero is a binary encoding scheme in which ones and zeros are represented by opposite and alternating high and low voltages, and where there is no return to a zero (reference) voltage between encoded bits.
NRZI	Non-return to Zero Inverted is a binary encoding scheme that inverts the signal on a "1" and leaves the signal unchanged for a "0".

N

NSN(s)	AMNET's Network Switching Node(s)
Nucleus 7000 or N7000	AMNET's series of networking products which includes the N7200, N7400, N7500, and N7900
Nucleus 7200 or N7200	AMNET's Multiprotocol Switching PAD
Nucleus 7400 or N7400	AMNET's Network Switch (MS-DOS™ based)
Nucleus 7500 or N7500	AMNET's Network Switch (UNIX™ based)
Nucleus 7900 or N7900	AMNET's Network Manager (UNIX™ based)
NUI	Network User Identification accommodates user ID, billing, and on-line facilities registration.

O

On-line	A connection to a computer enabling data to pass to or from the computer with human intervention.
Operating System	The software of a computer that controls the management and execution of programs, typically handling the functions of input/output control, resource sharing, and data management.
OSI	Open Systems Interconnection is a logical structure for the operations standardized within the ISO. It is a seven layer network architecture being used for the definition of network protocol standards to enable any OSI device to communicate with any other OSI compatible computer or device.

P

Packet	A group of bits consisting of data and control information that is transmitted as a composite whole over a physical link in a specific format, as the subset of a larger message.
Packet Switching	A data transmission process using addressed packets in which a channel is occupied only for the duration of transmission of the packet. Data and user information is segmented and routed in discrete data envelopes.
PAD	Packet Assembler/Disassembler is a network interface device that allows multiple asynchronous and/or synchronous terminals or host computer ports to interface to a packet switching network.
PDN	Public Data Network is a tariffed packet sw[illegible]g data carrier.
PIA	Parallel Interface Adapter is a high-speed interconnection between a NSN to a locally attached N7900 Network Manager.
Port	It is the entry or exit point from a node. Standard LPs and T-1/E-1 LPs have four (4) ports.
PPM	Packet Processing Module is AMNET's operating system resident in the N7400 and N7500.
PPS	Packet Per Second is a measurement of the rate or speed at which packets flow through the network.
Private Network	This is a network established and operated by a private organization or corporation confined to use by one customer.
Protocol	An agreement on format, meaning, and relative timing of information exchanged between two (2) communication devices.
PU	It is the abbreviation for Physical Unit and is a term used in the SNA environment to identify a printer, terminal or PC address.

P

Public Network It is a network operated by common carriers or telecommunications administrations for the purpose of providing circuit switched, packet switched, and leased line circuits to the public.

PVC Permanent Virtual Circuit is a logical connection between two (2) subscribers that is analogous to a leased line but supporting a single user.

Q

Q-Bit This is the Qualifier Bit in an X.25 packet that allows the DTE to indicate that it wishes to transmit data on more than one level.

QLLC Qualified Logical Link Control

R

RAM Random Access Memory

REJ Reject message tells the station to retransmit all information frames beginning with the response frame sequence number, which follows the REJ message.

RJE Devices Remote Job Entry devices are usually a card reader, printer, and terminal.

RNR Receiver Not Ready is used by a station to indicate a temporary inability to accept additional information frames, and to acknowledge previously received information frames.

Routing The path taken by packets as they traverse the network from source node to destination node.

Routing Tables Tables built either statically or dynamically updated by each node to define its neighboring nodes and to maintain current routing information.

R

RPOA	Registered Private Operating Agent permits the use of one or more networks to route a call to its destination. If more than one network is chosen, then the extended format is required.
RR	Receiver Ready is a message to indicate that it is ready to receive an information frame, and to acknowledge all previously received information frames.
RS-232-C	An EIA specified physical interface with associated electrical signaling between DCE and DTE devices.
RS-422-A	An EIA specification of the electrical characteristics of a balanced voltage digital interface circuit.
RS-423-A	An EIA specification of the electrical characteristics of a unbalanced voltage digital interface circuit.
RS-449	General-purpose 37-position and 9-position interface for DTE and DCE employing serial binary data interchange.
RS-530	High-speed 25-position interface for DTE and DCE employing serial binary data interchange; developed to serve as a complement to RS-232-D for data rates above 20 Kbps.

S

SDLC	Synchronous Data Link Control is IBM's version of CCITT's HDLC bit oriented protocol.
Shared Memory	Read/write access to this memory can come in from multiple sources and out to multiple sources. Also referred to as dual-ported memory or multi-ported memory.
SNA	IBM's Systems Network Architecture
Subscriber	A user of the devices accessing the network.

S

SVC	Switched Virtual Circuit is a temporary logical association between two (2) subscribers connected to a network, analogous to connection by a dial-up line.
SYNC	Synchronous transmission is a process where the information and control characters are transmitted at even intervals to preserve continuity (synchronization) within a data communications system.
System	A logical collection of computers, peripherals, software, service routines, terminals, and end users.

T

TCP	Transmission Control Protocol normally associated with IP, an internetworking software set originated on the DOD's network.
Throughput Class	The aggregate usable data rate for a particular port. Also spelled "thruput".
Topology	This refers to the overall configuration of a network such as star, ring, or mesh.
TPAD	The Terminal-to-Host PAD
Traffic	The volume and intensity of transmitted signals on a communications channel.
Transport Layer	Layer 4 of the OSI mode responsible for the end-to-end control of transmitted data and the optimized use of network resources.
Transparent Mode	In BSC data transmission, the suppression of recognition of control characters to allow transmission of raw binary data without fear of misinterpretation.
Trunk Protocol	A protocol used to handle the trunk line. Refer to LAPB.

V

V.24	A CCITT Recommendation that defines the function of the interchange circuits between a DTE and a DCE.
V.28	CCITT standard for electrical characteristics of interchange circuits.
V.35	A CCITT Recommendation for governing data transmission at 48 Kbps over 60 KHz to 108 KHz group band circuits.
VLSI	Very Large Scale Integration
Virtual Circuit	In packet switching, network facilities that give the appearance to the user of an actual end-to-end circuit where actual transmission facilities are shared by many users simultaneously.
VT100	A designation for an asynchronous DEC terminal.

X

X.3	A CCITT Recommendation that defines the service provided by a PAD to connect a character mode terminal to an X.25 network.
X.25	A CCITT Recommendation that specifies the interface between data terminal equipment and data switching equipment operating in the packet mode on public data networks.
X.28	A CCITT Recommendation that defines the user interface to the X.3 services.
X.29	A CCITT Recommendation that defines the use of X.25 packets to carry data between a terminal and a host.
X.25 Port	The physical connection in the network that provides the data paths between nodes.

X

X.121 — A CCITT Recommendation defining international addressing conventions applied to DTEs connected to public data networks.

XPAD — A transparent PAD that can be used to network any SNA device that can be networked using any of the other SNA PAD types. It can be considered a fully transparent PAD for synchronous protocols.

AMNET

1881 Worcester Road
Framingham, MA 01701 U.S.A.
(800) 846-2001 (U.S.A. only)
Phone: (508) 879-6306
FAX: (508) 872-8136