ORNL/TM-10177

ornl

OAK RIDGE NATIONAL LABORATORY

MARTIN MARIETTA

An Evaluation of Ethernet Connectivity via Broadband Bridge

Robert W. Hayes
John T. Farmer, Jr.

OPERATED BY
MARTIN MARIETTA ENERGY SYSTEMS, INC.
FOR THE UNITED STATES
DEPARTMENT OF ENERGY

Printed in the United States of America. Available from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Road, Springfield, Virginia 22161
NTIS price codes—Printed Copy: A03; Microfiche A01

ORNL/TM-10177
Dist. Category UC-37
Instruments

Instrumentation and Controls Division

AN EVALUATION OF ETHERNET CONNECTIVITY VIA BROADBAND BRIDGE

Robert W. Hayes
John T. Farmer, Jr.*

Publication Date: November 1987

*Midwest Technical, Inc.

Prepared by the
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, Tennessee 37831
operated by
MARTIN MARIETTA ENERGY SYSTEMS, INC.
for the
U.S. DEPARTMENT OF ENERGY
under Contract No. DE-AC05-84OR21400

CONTENTS

CONTENTS (continued)

LIST OF FIGURES

ABSTRACT

This report presents the results of an evaluation project to examine the performance characteristics of the Bridge Communications, Inc., Ethernet-Broadband bridge. This bridge is under consideration for use on the Oak Ridge National Laboratory/interplant broadband network to provide interconnectivity among various local Ethernet segments installed throughout the facility.

Evaluation of the bridge is based on the needs of plant-wide communication via broadband technology. Tests of performance and reliability demonstrated that the IB/1 Bridge offers adequate throughput for this application with excellent reliability and durability.

1. INTRODUCTION

1.1 FOREWORD

Presented here are the results of an evaluation project to examine the performance characteristics of the Bridge Communications, Inc., Ethernet-Broadband bridge. This bridge is under consideration for use on the Oak Ridge National Laboratory (ORNL)/interplant broadband network (ORNL/IPBB) to provide interconnectivity between various local Ethernet segments installed throughout the facility. Evaluation of the bridge is based on the needs of plant-wide communication via broadband technology, and a series of tests were performed on four Ethernet-Broadband bridges over a period of six months. Longer term performance is presently being observed but is not within the scope of this study.

1.2 OBJECTIVES

The following objectives were established as goals for this evaluation:

- Examine the Bridge Communications, Inc., Ethernet-Broadband bridge for potential use on the ORNL/IPBB.
- Determine performance characteristics of the bridge from both an Ethernet and an rf viewpoint.
- Determine limitations of bridges being tested.
- Determine limiting factors that affect total system performance (e.g., maximum throughput of backbone network, bridge throughput, etc.).

The four bridges were evaluated based on the criteria listed below.

- Protocol transparent operation.
- Ability to operate over long broadband cable distances.
- Ability to operate reliably under normal broadband signal variations and to perform under mild to moderate broadband disturbances.
- Provide high throughput and packet filtering rates.
- Provide good network management and security features.

2. METHODOLOGY EXAMINATION

An important part of the bridge testing was an examination of the available features of the bridges and the backbone broadband network. Section 2.1 describes how these issues are addressed.

The bridge evaluation produced two types of information concerning the bridges under test: physical attributes and overall network performance. The first area addresses such issues as the waveform characteristics of the broadband modem used, the ease of network management, and the actual construction and maintainability of the bridge. The second addresses the actual performance of the bridges as would be seen by users of the network. Sections 2.2 and 2.3 describe the methods used to address these points.

2.1 FEATURES TESTING

The bridge was examined and tested to determine its salient characteristics and its capabilities, which include capabilities as tested in this evaluation and described in the documentation supplied. Future and anticipated features are omitted, as well as any feature not implemented or in operation at the time of testing.

2.2 BACKBONE NETWORK TESTING

The Ethernet bridge examined uses one pair of standard CATV channels on the broadband network. (A channel pair consists of one forward and one reverse channel.) The protocol used was examined for its affect on the throughput of the Ethernet-to-Ethernet transmission, its ability to handle a range of network loading, and its ability to perform well over long physical cable distances.

2.3 PHYSICAL ATTRIBUTES TESTING

The physical attributes of the bridge (maintainability, broadband characteristics, and setup and management capabilities) were examined and compared with current and future requirements. These characteristics are of great importance due to the high reliability requirements for the network.

2.3.1 Maintainability

The bridge was evaluated for its maintainability both while it is in use and while it is removed for repair. Evaluation was based on ease of installation and repair, diagnostic capability, and bridge reliability.

2.3.2 Rf Characteristics

The waveform characteristics, bandwidth requirements, susceptibility to noise, and adaptability to a changing network configuration were evaluated by laboratory testing and field measurement of the devices. Due to the extensive cable network, any device connected to the broadband network must be durable enough to operate in less than optimal conditions. Therefore, this area of the evaluation is of utmost concern, since many vendor promises cannot be realized on a network as vast as the ORNL/IPBB.

2.3.3 Bridge Setup and Management

In examining the ease of setting up each bridge for use in the network was any physical and/or software network reconfiguration that might be required.

Vendor procedures for setting up a new bridge[1] were reviewed and followed. Evaluation of the bridge was based on functions that could be performed at the bridge and remotely through the network[2] and on the flexibility and capability for network management.

2.4 PERFORMANCE TESTING

Two throughput parameters are of interest for network bridge devices: the maximum packet forwarding rate and the actual packet filtering rate.[3] These two parameters define the maximum network load that the bridge can support without contributing to a network bottleneck.

The maximum packet forwarding rate specifies the total amount of data that can be passed through the bridge under optimal conditions. The bridge performs local packet traffic filtering; therefore, not all data packets seen on the Ethernet are forwarded. Although the local traffic is not forwarded onto the broadband, each packet, regardless of its destination, must be examined by the bridge to determine if it is to be forwarded.

The packet filtering rate must be adequate to handle traffic bursts without losing any packets. Packets that are lost must be retransmitted by the sending node, causing undue delay and additional overhead.

With the equipment on hand, there were three series of tests that could help determine the packet forwarding and filtering rates. These three tests would allow measurement of the raw packet forwarding performance, task-to-task communication performance using a higher level protocol, and performance of a "command line" level file transfer.

The packet throughput testing performed on the bridges was done with only the two test devices on the network operated in an active sender, passive receiver mode. By testing in this manner, maximum bridge performance in an optimum environment could be obtained.

2.4.1 Raw Packet Throughput

Each bridge was tested for the maximum number of packets it can pick up and forward to the backbone network. This forwarding ability was tested in two configurations using the ORNL/IPBB. The first configuration (Fig. 2.1) involved installation of the two bridges on two separate Ethernets in the same laboratory. This laboratory setup was located close to the headend to minimize distance-dependent interactions. In the second configuration, one bridge was located near the headend, and the second was relocated to a distant site at the Oak Ridge Y-12 Plant (Fig. 2.2).

Two Ethernet analyzers were used, one connected to each Ethernet. One analyzer generated data packets at varying sizes and rates, while the other was used to monitor the number of packets arriving at the second Ethernet. Testing was run independently in both directions in an attempt to minimize any impacts of differing analyzer throughput, send/receive rates, and other capabilities.

2.4.2 DECnet Testing

Bridge performance was also tested, using application-level protocols such as DECnet. Performance was measured for node-to-node file transfers and for task-to-task communication. Execution of the test series was in both directions between the nodes in an attempt to account for differences in node performance.

Software consisting of a master/slave pair of codes was developed and executed on different machines generated load conditions and measured performance characteristics. The master code software opens a DECnet channel and requests activation of the slave code. When communication is established, the master code transmits data blocks to the slave code, which receives them and returns acknowledgment messages to the master. The master code keeps statistics on the time required to open the channel, transmit a set of data blocks, and close the channel. Appendix A contains the pseudo-code for the master and slave codes.

Testing was done with data block sizes from 1 to 4096 bytes, with particular attention to block sizes of 40 and 512 bytes. The 40-byte block was chosen to examine throughput for small blocks wherever software overhead and header overhead are high. This measurement should demonstrate approximate performance under interactive traffic such as terminal-to-host connections. Keyboard traffic will be transmitted probably one or two characters at a time but at very low rates. The return data from the computer to the CRT screen will be in blocks varying from a few characters to several lines. A good average for this performance is about 40 characters (or 40 bytes), which means

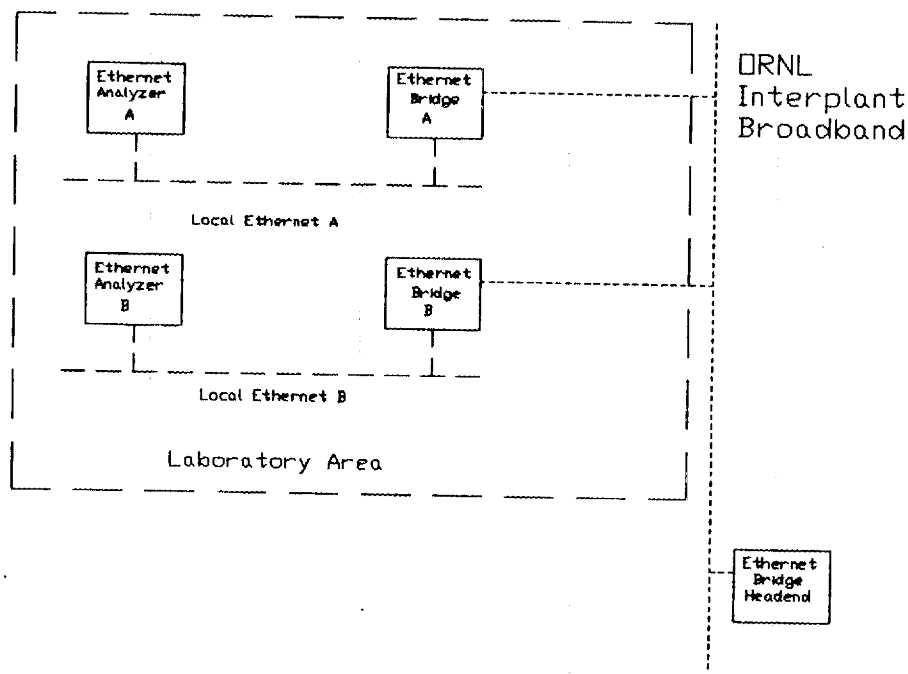

Fig. 2.1. First configuration for data throughput tests.

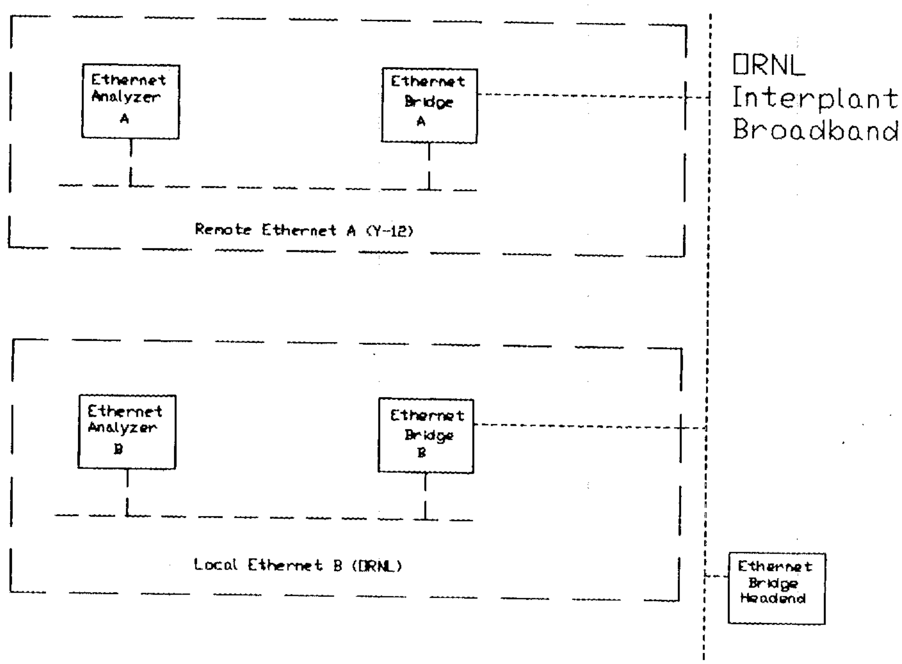

Fig. 2.2. Second configuration for data throughput tests.

that the performance measurements made at the 40-byte-block size should be representative of the performance seen during normal terminal screen output.

The 512-byte data-block size was chosen as representative of most task-to-task communications and system-to-system file transfers. This size also fits into the standard DECnet default maximum data buffer size of 576 bytes.

2.4.2.1 Test configuration. Application-level testing was performed using two network configurations. The first used a single Ethernet, and the other used the bridges to connect the two Ethernets. This test characterizes the bridges' performance in relationship to a standard network. (See Sect. 2.5 for more specific information on test procedures and setups.)

2.4.2.1.1 Local Ethernet connection. The test series was first executed between two DECnet nodes located on the same Ethernet, which establishes the performance baseline for comparison purposes.

2.4.2.1.2 Bridge connection. The second test series was executed between two Ethernets interconnected by the bridge device. The only path between the source and the destination test partners was through the bridge under test.

2.4.2.2 Test types. Two application-level tests were developed: the first measures throughput capability between two cooperating programs, and the second measures the throughput available to a user transferring data at a command level interface.

2.4.2.2.1 Task-to-task communication. Two tasks were set up in a master/slave configuration, where the master task invokes the remote task and passes to it the size of buffers it will be sending. The master then sends a known number of test buffers to the slave for several iterations and reports the time used for each iteration. The slave process, when activated, waits for the "buffer size" message. The slave then accepts this message and accepts each subsequent data buffer as it arrives.

Refer to Appendix A for the pseudo-code used for the master and slave processes.

2.4.2.2.2 File transfer. A set of data files was created and copied using standard system "copy" commands. The data files range in size from zero bytes to a large number of bytes (on the order of 280,000 bytes). Each file was copied from the source node to the destination node several times and the time used for a single file transfer was calculated.

Refer to Appendix B for the skeleton command procedure used for this test.

2.5 TEST PROCEDURES AND CONFIGURATIONS

Each test and evaluation configuration and procedure was documented as described below. All command files and programs are listed in Appendix B.

2.5.1 Physical Maintainability

Each bridge was examined for maintainability features and utilization of standard devices. The evaluation considered the difficulties encountered in checking and replacing components within the bridge and in diagnostic features available in the device.

2.5.2 Rf Performance Measurement

The rf performance of the bridges was a significant issue due to the long physical length of the ORNL/IPBB. This large size makes controlling the gain, loss, tilt, flatness, and noise ratios more difficult to control. Each bridge was installed on a test broadband network to examine the modem characteristics and to determine if it was appropriate for use on the active broadband system. The modem output signals were examined for signal level, uniformity, out-of-band signals, spurious noise, and noise floor. The maximum and minimum signal levels at which the bridges would operate reliably were determined by varying the test network loop loss in both forward and reverse paths.

After the modems were installed on the active broadband network, repeat measurements were made to verify performance on a large network. Also, the bridges were observed for susceptibility to noise and to signal variations.

All observations of signal quality, signal level, and other characteristics were made utilizing an rf spectrum analyzer. Signal level adjustments were made utilizing rf attenuator pads of various values.

2.5.3 Performance Testing

The following sections describe the procedures, configurations, and equipment used to perform the various packet throughput measurements and calculations.

2.5.3.1 Packet throughput. The equipment needed for performing the packet throughput testing consists of

- Two Ethernet analyzers capable of generating and receiving packets of arbitrary size, type, source address, and destination address.

- Two unused Ethernet cable plants (suggested: DEC Delnis or equivalent).
- A properly installed backbone network.

Figure 2.3 illustrates the test equipment configuration.

The two Ethernets were attached to the bridges with one analyzer connected to each Ethernet. One analyzer (denoted as "B") was set up to receive, error check, and count all test packets that appear on its local Ethernet. The other analyzer (denoted as "A") was set up to generate 10,000 packets of a specified size, using a user-specified delay period between each packet sent. The operator then used analyzer A to place a burst of test packets on its Ethernet. The time required for the packets to arrive at analyzer B was measured. This procedure was done using several different packet delay times (usually 0- and 1-ms delays are adequate) and with packet sizes [excluding preamble and FCS (frame check sequence)] of 64 to 1512 bytes. The following values were recorded for each pass:

- Time required to transmit all packets in the test pass,
- Number of packets placed on the source Ethernet,
- Number of packets monitored by the receiver analyzer,
- Number of packets examined by source Ethernet bridge interface,
- Number of packets forwarded onto the backbone network, and
- Number of packets transferred from backbone network to destination Ethernet.

Following is a list of the calculations performed using the raw data recorded from the packet throughput performance tests.

1. Supplied packet rate was calculated by dividing the number of packets sent from the source analyzer by the amount of time used. This rate defines the network load being presented to the bridge under test.

   (num_packet_sent/time_used)

2. Packet pickup rate (Ethernet to IPBB) was calculated by dividing the number of packets forwarded to the broadband channel by the amount of time used. This rate defines the actual capability of a single bridge to accept and forward packets from the attached Ethernet.

   (num_packet_forwarded/time_used)

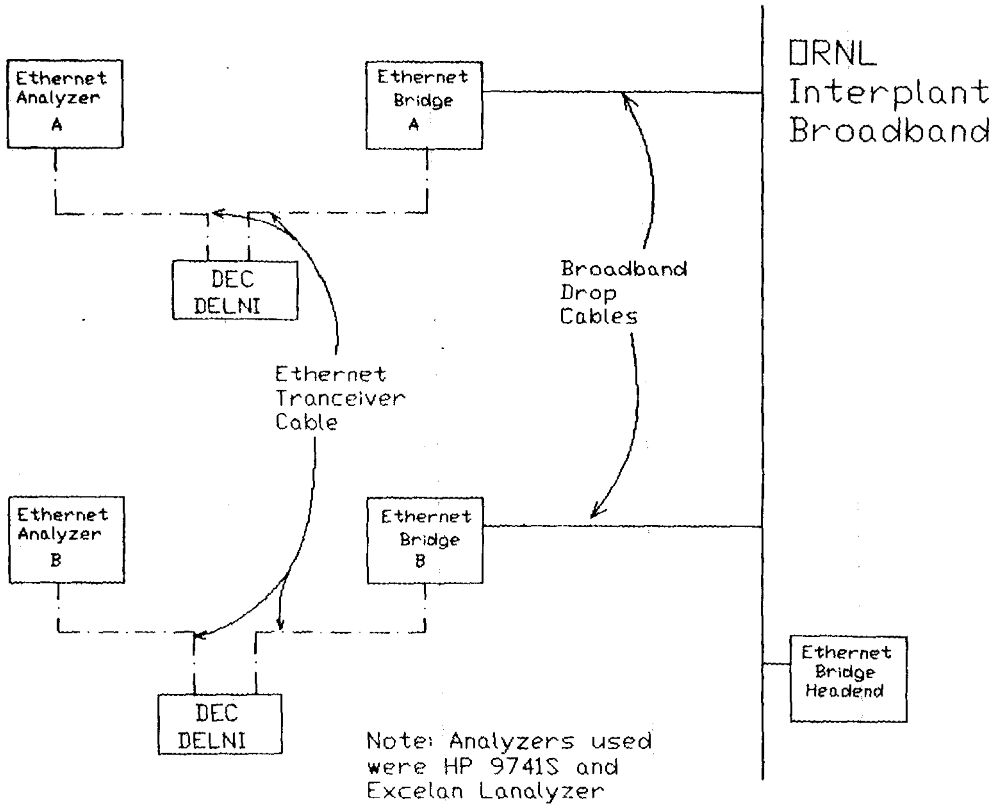

Fig. 2.3. Bridge test equipment and connections.

3. The error rate from Ethernet pickup to receiving analyzer was calculated by subtracting the number of packets seen at the receiving analyzer from the number of packets forwarded onto the broadband. This value is then divided by the number of packets forwarded to the broadband.

   ((num_packet_forwarded - num_packet_received) / num_packet_forwarded)

4. Total packet throughput was calculated by dividing the number of packets seen at the receiving analyzer by the amount of time used.

   (num_packet_received/time_used)

5. Megabit per second throughput was calculated by converting the total size to bits, adding the total number of bits used for the preamble and the FCS, multiplying this value by the number of packets per second, and then dividing by the value 2^{20} (converts from bits per second to megabits per second).

$(((packet\_size*8) + preamble\_bit\_size + FCS\_bit\_size) * num\_packets\_per\_second) / (2^{20})$

2.5.3.2 DECnet performance. To perform the application level testing with an Ethernet bridge, the following was used:

- Two DEC VAX computers equipped with Ethernet interfaces and DECnet licenses,
- Two properly installed Ethernets (DEC Delni's or equivalent), and
- A properly installed backbone network.

Test programs and data files were installed on each VAX node and were verified to operate correctly on a local basis. The two VAX computers were first connected to the same Ethernet for performing the baseline testing. The two types of tests were then executed using one node (denoted as node "A") as the source; the tests were repeated using the other node (denoted as node "B") as the source (Fig. 2.4a).

Node B was then removed from the single Ethernet and connected to the second Ethernet. The bridges were installed to interconnect the two Ethernets. The two test types were then repeated, using node A and then node B as the source (Fig. 2.4b).

2.5.3.2.1 Task to task. This test was designed to measure the throughput for applications that must communicate buffers of information between two or more processes through the VMS/DECnet interface. A command file was set up to execute the MASTER program for data blocks of size 1 to 4096 bytes. This command file was submitted for overnight execution on the host or originator system.

2.5.3.2.2 File transfer. This test was configured to measure the throughput that can be expected using high-level file manipulation commands (Copy, Delete, Directory, etc.).

A scratch area located on a non-system disk drive was set up on each system involved in the testing. An ASCII text file consisting of varying length text records was used as the transfer data, and several versions ranging in size from 0 bytes to ~280,000 bytes were constructed. A command file was written to repeatedly copy the various files between two systems and was submitted for overnight execution. (Refer to Appendix B for a sample command file.)

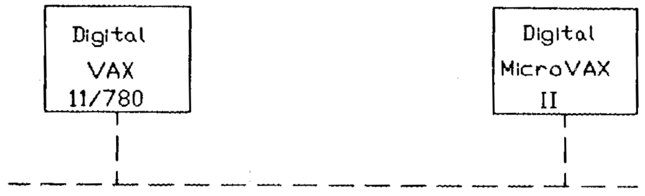

(a) Local Ethernet connection

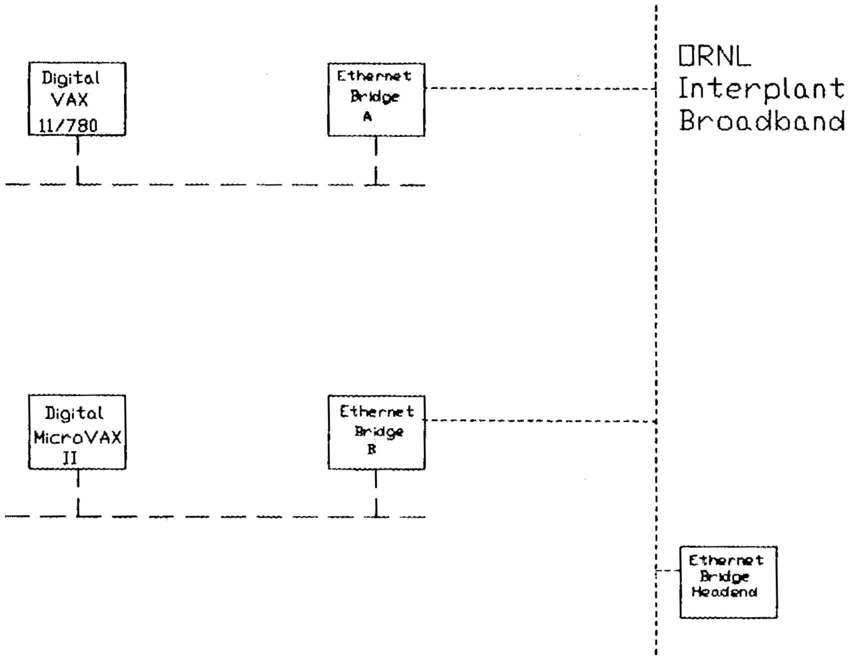

(b) Connections via bridge

Fig. 2.4. Throughput tests via local Ethernet and bridge connections.

3. EVALUATION

The Ethernet bridge devices were tested and evaluated as described above. The bridges were subjected to laboratory tests followed by field testing by the evaluation team and users of Energy Systems unclassified computer networks. The results of these tests are presented below.

3.1 FEATURES EVALUATION

The following table is a summary of features provided on the units tested. A brief comment on the quality of implementation provided on the particular bridge is listed in the appropriate column. A detailed discussion of the features would be excessive and beyond the scope of this report. Additional information may be obtained by contacting one of the authors or the vendor.

Feature	Bridge
Protocol independence	Excellent
Ability to operate control mode without interruption of service	Yes
Local traffic filtering	Good
Selective filtering by address	Good
Selective filtering by packet type	Good
Password security	Good
Settable password	Yes
Statistical data collection	Good
Network management functions	Fair
Central network management features	Yes
LED status indicators	Good
Power fail recovery	Good
Coordinated time clocks	Yes
Adaptive transmit signal level adjustment	Yes
Frequency agile	Yes

3.2 BACKBONE NETWORK EVALUATION

The bridge operates using the ORNL/IPBB as the means of interconnecting multiple Ethernets. It uses one 6-MHz channel for its forward signal and a second 6-MHz channel for its reverse signal. Table 3.1 is a list of backbone characteristics.

Table 3.1. Bridge IB/1 characteristics

Channel bandwidth	1 -forward 1 -reverse
Basic data rate (megabits/second)	5
Channels available (forward/reverse)	4A'/ R 5' / S 6' / T FM1'/U FM2'/V FM3'/W
Data encoding Method	Duo-binary PSK
Type of protocol	Slotted CSMA/CD

The Bridge IB/1 uses a protocol that is for all intents and purposes identical to standard CSMA/CD (Carrier Sense Multiple Access/Collision Detect)[4] Ethernet. The major differences arise because of the different speed of the network (5 Mbits/s) and the long length of the associated cable plant. The Ethernet standard restricts the maximum cable length to 1000 m, which allows a 5-μs slot time for each packet at the 10-Mbit/s data rate. Due to the long distance that the signals must travel in a facility broadband network and the lower data rate used by the IB/1, the slot time used for a packet must be increased and the minimum packet size increased. For the ORNL/IPBB, the slot time required is 408 μs (based on the assumption that the last bridge is 15 miles from headend). Each packet sent on the broadband is padded to the minimum packet length (presently 256 bytes) and placed on the IPBB, where it is read and checked by each IB/1 on the channel. If the destination address of the message matches an address on an IB/1's list, that IB/1 then forwards the message to its local Ethernet. Otherwise the packet is ignored.

3.3 PHYSICAL ATTRIBUTES EVALUATION

Bridges selected for use on the ORNL/IPBB, in addition to providing adequate network throughput, must be maintainable, work within the ORNL/IPBB, and provide capabilities for network management.

Access to external and internal components of the IB/1 is easily accomplished. The local floppy disk drive is accessible from the front of the IB/1. The cover is removed by lifting the back edge of the top and popping two spring clips concealed in the unit's sides. After the cover has been removed, all of the internal components may be reached, and there is ample room around each one.

The IB/1 provides six LEDs on the front panel for system status monitoring. Their labels and functions are

- Network activity – flashes when packets are observed on the broadband channel.
- Packet received – flashes when the IB/1 picks up a packet on the local Ethernet.
- Packet forwarded – flashes when a packet is forwarded from the local Ethernet to the IPBB or when a packet is forwarded from the IPBB to the local Ethernet.
- Boot state – turns on when the IB/1 is rebooted or turned on. Stays on until the booting process is completed and the unit is ready to operate.
- Self-test – is on while the IB/1 executes a series of internal self-check diagnostics.
- Power – is on whenever power is supplied to the unit and the power switch is on.

The Bridge IB/1 is built in a modular fashion that appears to allow field replacement of almost all of its components. The IB/1 uses the Multibus (IEEE Std 796) as its basic backplane interconnection system. The unit consists of three multibus cards in a seven slot cardcage, one half-height 5.25-in. DS/DD floppy diskette drive, a sealed power supply, and a muffin cooling fan. The three multibus cards are the Ethernet interface, broadband interface, and a single-card, 68000-based microcomputer. The cards, disk drive, power supply, and fan all appear to be easily replaceable. The broadband interface connects to an external RFM/5 Bridge broadband modem unit through an 18-in. cable terminated with 37-pin D-shell connectors. As such, the broadband modem can be replaced without opening the IB/1.

To date, we have had little experience with diagnosing and repairing the Bridge IB/1 because none of the units has failed. Of the four that were purchased and installed, there have been no hardware failures either in startup or in service. The vendor has provided one firmware update and several software updates to accommodate the large IPBB cable plant.

The IB/1 bridge has a 5.5-MHz-wide signal and operates on a remodulator rather than on a translator. This means that the headend fully receives the data packet and remodulates it back out in the forward direction, giving the outgoing signal a minimum signal-to-noise ratio and a uniform signal level. Also, other devices can be installed directly adjacent to the bridge channel. The bridge has a +/- 12-dB dynamic range with automatic signal level adjustment of +/- 14 dB. It is tolerant up to 26 dB of additional attenuation in the reverse direction, although the forward direction can only tolerate 12 dB of extra attenuation. The bridge has not shown any sensitivity to ambient noise, and it was installed easily in Bldg. 9201-2.

3.4 BRIDGE SETUP AND MANAGEMENT

The installation of the IB/1 Bridge involves only connecting it to the broadband and to the Ethernet. The IB/1 automatically adjusts its transmitter signal level to within tolerances specified by thumbwheel switches on the headend remodulator.

4. PERFORMANCE TESTING

The bridges were tested for performance using a variety of tests in an effort to determine both their optimum and real-world performance.

4.1 RAW PACKET THROUGHPUT

Table 4.1 shows the calculated total packet forwarding rate, error rate, and megabit/s throughput of the Bridge IB/1. As can be seen by examining the table, the worst forwarding rate exhibited by the Bridge IB/1 is for the very largest allowable Ethernet packets (size 1500 to 1514 bytes). This value is in the 380- to 400-packet/s range. The forwarding throughput steadily increases as the packet size is decreased until the packet size reaches 512 bytes. For packets smaller than this, throughput jumps to the 1600- to 2000-packet/s range. Traffic on Energy Systems Ethernets is typically 90% packets smaller than 512 bytes in size; therefore, the available throughput of the IB/1 should be in the range 1500 and up packets/s.

The error rate shown is based on the number of packets that were lost in the forwarding process, which was tested during daytime hours using the actual ORNL/IPBB. Under these conditions the error rate did not exceed 0.36%.

The throughput expressed as megabits per second shows that at the typical packet sizes expected, the IB/1 is only using approximately 20% of the bandwidth available in its broadband channel.

Figure 4.1 is a graph of the supplied packet and the packet pickup rate as plotted against the packet size. The packet supply rate is the number of packets being placed on the source Ethernet by the generating analyzer. This value is then compared with the packet pickup rate to determine the effect on the bridge forwarding capability as it relates to Ethernet traffic density. As can seen in Fig. 4.1, the packet pickup rate of Bridge IB/1 tracks the traffic density closely. Bear in mind that in this scenario, all the traffic appearing on the local Ethernet must be forwarded onto the backbone IPBB. Bridge claims that the IB/1 can examine and filter up to 7000 packets/s and forward up to 2500 packets/s. If these numbers are for a small broadband cable plant, our measured forwarding and pickup peak of 1900 packets/s is in fair agreement with their claims. Without additional equipment to generate local traffic and forwarding traffic, we could not prove or disprove the filtering and examining claim of 7000 packets/s. However, it would not be unreasonable, in light of the other data, for this to be a valid upper bound.

Figure 4.2 is a graph of the total throughput rate plotted as packets per second against packet size. As can be seen by examining the graph, the Bridge IB/1 packet throughput is directly related to the packet

Table 4.1. Bridge IB/1 throughput data

Packet size (bytes)	Total packet forwarding rate (packets/s)	Error rate (% transmitted)	Throughput (Mbits/s)
64	1884.5	0.05	0.92
80	1864.9	0.05	1.14
128	1935.1	0.04	1.89
200	1895.7	0.09	2.89
210	1811.1	0.00	2.90
220	1775.4	0.08	2.98
230	1707.2	0.05	2.99
240	1672.1	0.04	3.06
256	1627.4	0.05	3.18
512	1027.4	0.01	4.01
1024	563.2	0.03	4.40
1400	420.7	0.03	4.49
1500	393.0	0.12	4.50
1510	393.0	0.15	4.53
1512	392.6	0.07	4.53
1514	385.8	0.36	4.46

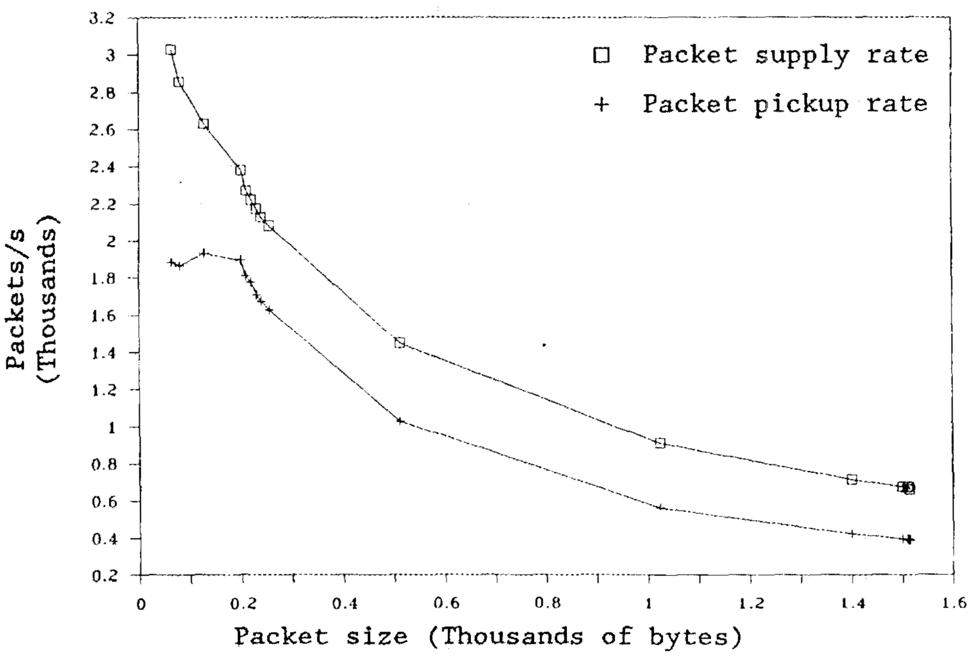

Fig. 4.1. Packet pickup and forwarding rates.

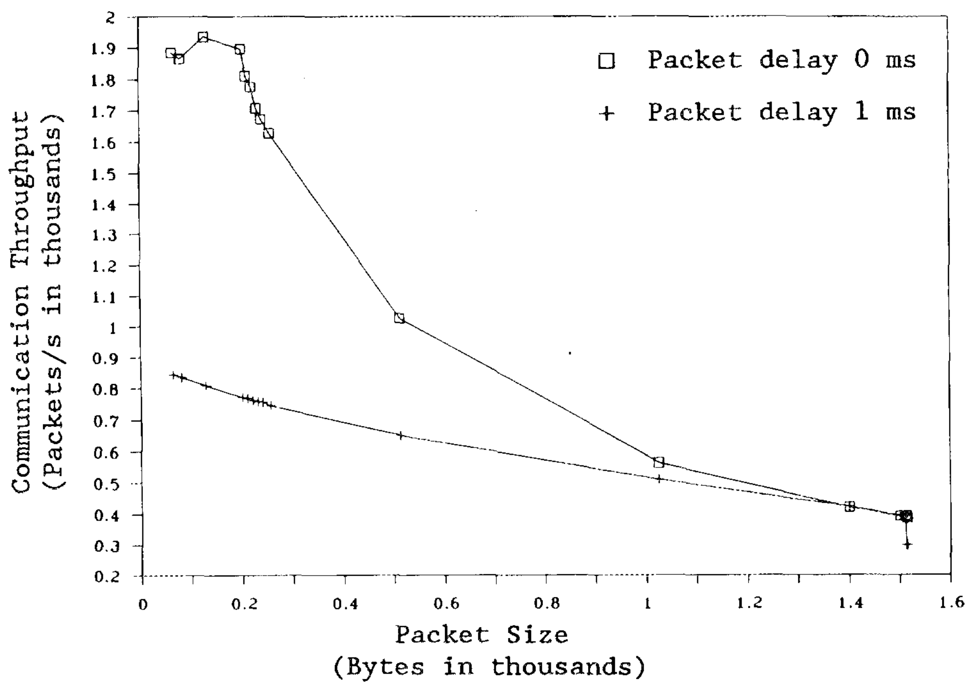

Fig. 4.2. Bridge IB/1 performance: total throughput rate plotted as packets/s against packet size.

size. The best performance for continuous forwarding demands occurs with packet sizes of 256 bytes and smaller. This peaks in the 1800- to 1900-packets/s range (see upper line on graph). The lower line on Fig. 4.2 represents a forwarding load on the bridge that is still high; however, the packets are not arriving continuously. In this situation, the IB/1 is able to keep up with the data packet rate up to the largest packet size.

Figure 4.3 is a graph of the total throughput rate plotted as megabits per second against packet size. Examining this data, we see that the smaller packets have lower total data throughput due to the packet handling overhead. As the packet size is increased, allowing more time between packet handling, the microprocessor and the network interfaces can work more in parallel. This is evident in the fact that as the packet size is increased to the range of 1000 bytes and larger, the measured throughput approaches the maximum available through the 5 Mbit/s IPBB channel. This linearity can also be seen when a delay is inserted between packets; the throughput grows at a more linear rate and approaches the same limitations for very large packets.

A lower throughput for packets of size 1510 to 1518 bytes, the largest packets defined for Ethernet, appears to be caused by a bug in the

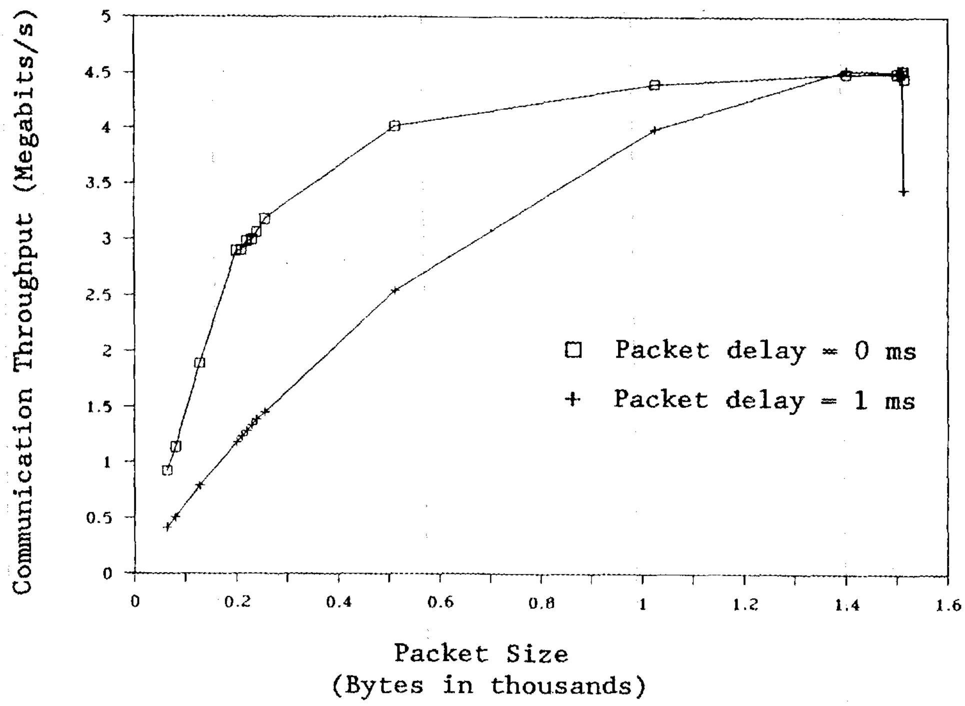

Fig. 4.3. Bridge IB/1 performance: total throughput rate plotted as megabits/s against packet size.

IB/1's internal memory management. This has been discussed with the vendor, who is investigating the problem.

4.2 DECNET TESTING

Testing of the Bridge IB/1 was conducted using the main I&C Division Ethernet and a test Ethernet set up for testing purposes. The three systems used were a VAX 11/780 (ICENET), a MicroVax II (ICMMIS), and a second MicroVax II (ICMPSD). (Acronyms in parenthesis are node names used to identify the machines.) Throughput testing of a local DECnet/Ethernet connection was done using only two I&C systems, ICENET and ICMMIS.

4.2.1 Local DECnet/Ethernet Performance

The executable files, data files, and command files were set up on the I&C VAX 11/780 (ICENET) and on a MicroVAX II (ICMMIS). The tests were executed in both directions to minimize any impact caused by the different types of systems and Ethernet controllers (DEUNA on the VAX 11/780 and DEQNA on the MicroVAX).

4.2.1.1 Task-to-task communication. Figure 4.4 is a graph of data throughput charted against data block size. As can be seen from the graph, the throughput of DECnet/Ethernet is almost linear in relationship to data block size. This linearity is due to the method Digital uses to package large data blocks for transmittal on behalf of a user process. A single Ethernet packet is used for a data block until the size of the data plus the DECnet overhead exceeds 1512 bytes (slightly under the maximum Ethernet packet size allowed). At this point, DECnet uses a second packet to contain data and a third packet with formatting information. As far as we have been able to determine, DECnet then adds additional data packets until the entire data record has been transferred.

4.2.1.2 File transfer. File transfer measurements were made, but all data values had very high scatter. It was clear that the Ethernet performance values were greatly affected by other factors, such as disk access, other system loads, and network statistics. Since file transfers occur in 512-byte blocks, the task-to-task data taken at 512-byte blocks is considered sufficient to indicate the overall performance expected for file transfers.

4.2.2 Bridge Performance

For testing the Bridge Communications IB/1, the executable files, data files, and command files were set up on the I&C VAX 11/780 (ICENET) located at ORNL Bldg. 3500 and on a MicroVAX II (ICMPSD) located at Y-12 Bldg. 9201-3. The tests were executed in both directions to minimize any impact caused by the different types of systems and Ethernet controllers (DEUNA on the VAX 11/780 and DEQNA on the MicroVAXs).

4.2.2.1 Task-to-task communication. Figure 4.5 is a graph of data throughput charted against data block size for the Bridge IB/1. The graph shows significantly less throughput than the raw bridge throughput. This reduction is due to the fact that the Ethernet interface in the VAX computer is only capable of 1.2 to 1.5 Mbits/s throughput. Also, the DECnet protocol requires acknowledge packets for the data packets sent. The time spent waiting for the acknowledgments reduces the total possible throughput.

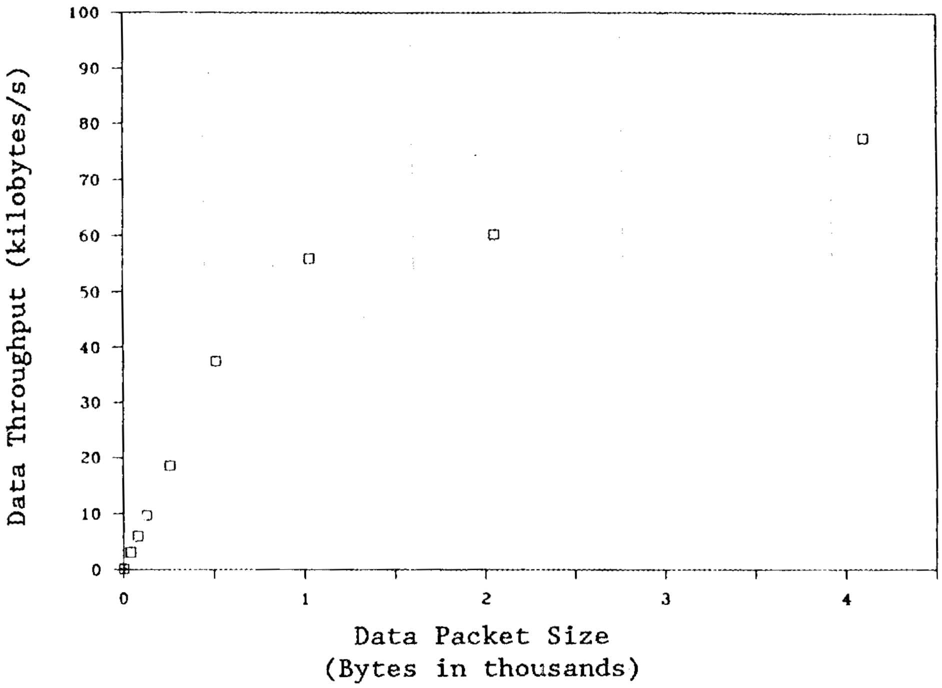

Fig. 4.4. DECnet/Ethernet data throughput charted against data block size for local Ethernet connection.

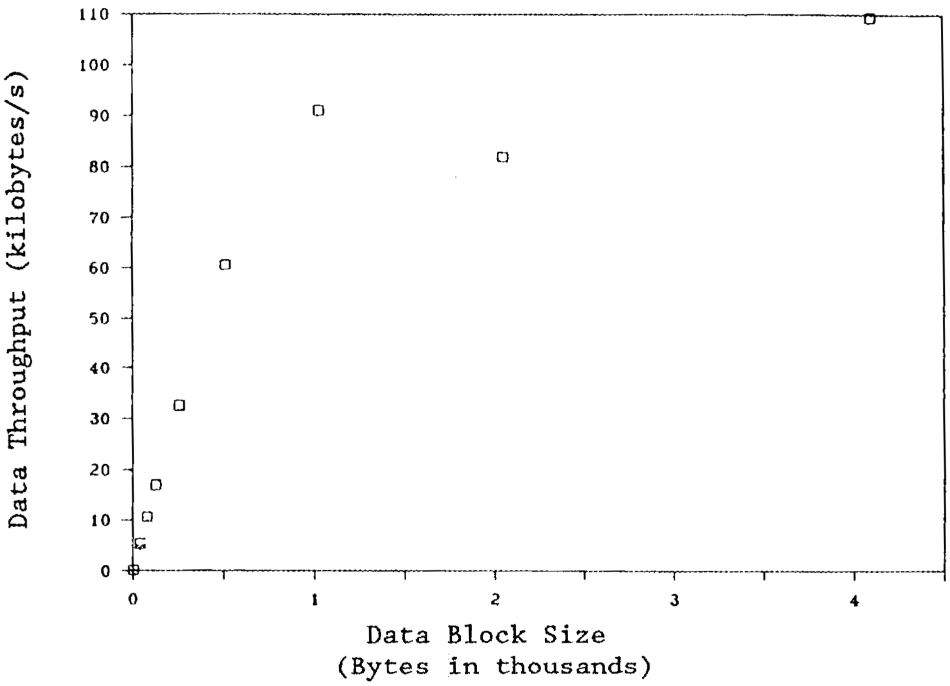

Fig. 4.5. DECnet/Ethernet data throughput charted against data block size for the IB/1 Bridge.

5. CONCLUSION AND RECOMMENDATION

Ethernet bridges were tested extensively for applications on the ORNL/IPBB. Tests of performance and reliability were performed, and a comparison is presented. The IB/1 Bridge offers adequate throughput with excellent reliability and durability, and it is therefore recommended that the Ethernet Bridge IB/1 be accepted for use.

APPENDIX A

PSEUDO-CODES AND COMMAND PROCEDURES

APPENDIX A. PSEUDO-CODES AND COMMAND PROCEDURES

MASTER PSEUDO-CODE

```
Get the number of repeats (K)
Get the number of inner loops (N)
Get the block size

Do K times
  Initialize timer
  Connect to remote task

  Send block size as connect argument data
  Read timer, store as OPEN_TIME
  Initialize timer

  Do N times
    Send data block
  End do

  Read timer, store as TRANSFER_TIME
  Initialize timer
  Disconnect from remote task
  Read timer, store as CLOSE_TIME
End do

Print timing information
Exit
```

SLAVE PSEUDO-CODE

```
Accept connection
Find block size from connect argument data

Do while connection open
   Receive block
End do

Close network access
Exit
```

The following is a description of the steps used to execute the task-to-task communication testing.

The symbol "SLAVETASK" is used in the master code as the logical file or device written to. It is set up to define where the receiving process is located. The syntax for defining this symbol is:

```
$def SLAVETASK "ICENET""SLAVEACCNT Password""::""task=slave"""
```

This specifies that the symbol "SLAVETASK" points to the command file SLAVE.COM located in account area SLAVEACCNT on node ICENET. On invocation by the master process, the network will search for and execute this command file. This command file will in turn run the executable image of the slave process.

The timing data generated by the master process are captured in the batch run log produced by VMS. To reduce variance due to different system loadings during the day, it is suggested that the tests be run as batch jobs overnight, saving the batch run log in a disk file for later analysis.

APPENDIX B

FILE TRANSFER SKELETON COMMAND PROCEDURE

APPENDIX B. FILE TRANSFER SKELETON COMMAND PROCEDURE

The command procedure is set up to transfer 15 data files from one system to a second system using DCL commands. The file is copied ten times in order to reduce variance from network traffic effects and from system loading. The system time is checked at the start of a test case, after the transfer has been completed and after the files have been successfully deleted from the target system.

```
$! command file to measure file transfer performance via Decnet
$! command line operations.  This will be done as an overnight
$! batch job.  This file "filexfer2.com" executes the tests
$! between the I & C VAX 11/780 (ICENET::) and the
$! I & C development MicroVax II (ICMPSD::)
$!
$! written by J. T. Farmer I & C 4/23/87
$!
$! We will execute 15 test cases with each case consisting of
$! 10 file transfers.
$!
$! T E S T  C A S E  # 1 File size = 565 blocks
$!
$ show system
$ show time
$copy drdl:[farmer.bridge_test]empty_file0.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 1
$copy drdl:[farmer.bridge_test]empty_file0.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 2
$copy drdl:[farmer.bridge_test]empty_file0.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 3
$copy drdl:[farmer.bridge_test]empty_file0.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 4
$copy drdl:[farmer.bridge_test]empty_file0.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 5
$copy drdl:[farmer.bridge_test]empty_file0.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 6
$copy drdl:[farmer.bridge_test]empty_file0.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 7
$copy drdl:[farmer.bridge_test]empty_file0.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 8
$copy drdl:[farmer.bridge_test]empty_file0.txt icmpsd::$disk2:-
```

```
[farmer.scratch_space]empty_file.txt
$! pass 9
$copy drdl:[farmer.bridge_test]empty_file0.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 10
$show time
$!
$delete icmpsd::$disk2:[farmer.scratch_space]empty_file.txt;*
$show time
$!
```

Each test case transfers the data file ten times and then deletes it. Test cases 2 through 14 are identical to test case 1 and test case 15. The file sizes are chosen to be approximately equal in increments from the largest (565 blocks, test case 1) to the smallest (0 blocks, test case 15).

```
$!
$! T E S T  C A S E  # 15  File size = 0 blocks
$!
$show system
$show time
$copy drdl:[farmer.bridge_test]empty_file14.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 1
$copy drdl:[farmer.bridge_test]empty_file14.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 2
$copy drdl:[farmer.bridge_test]empty_file14.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 3
$copy drdl:[farmer.bridge_test]empty_file14.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 4
$copy drdl:[farmer.bridge_test]empty_file14.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 5
$copy drdl:[farmer.bridge_test]empty_file14.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 6
$copy drdl:[farmer.bridge_test]empty_file14.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 7
$copy drdl:[farmer.bridge_test]empty_file14.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 8
$copy drdl:[farmer.bridge_test]empty_file14.txt icmpsd::$disk2:-
[farmer.scratch_space]empty_file.txt
$! pass 9
$copy drdl:[farmer.bridge_test]empty_file14.txt icmpsd::$disk2:-
```

```
[farmer.scratch_space]empty_file.txt
$! pass 10
$show time
$!
$ delete icmpsd::$disk2:[farmer.scratch_space]empty_file.txt;*
$show time
```

REFERENCES

1. Bridge CR/5 and RFM/5 Installation Guide, 09-0090-00, Bridge Communications, Inc., November 1986.

2. Bridge IB/1 User's Guide, software SW/IB version 10003,86-0073-00, Bridge Communications, Inc., August 1986.

3. A. S. Tannenbaum, Computer Networks, Prentice-Hall, Englewood Cliffs, N.J., 1981.

4. ANSI/IEEE Standard 802.3-1985, IEEE Standards for Local Area Networks: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications.

ORNL/TM-10177

Distribution Category
UC-37, Instruments

INTERNAL DISTRIBUTION

1. D. L. Aslinger
2. E. E. Branton
3-6. D. A. Clayton
7. J. E. Cleaves
8. J. T. Cleveland
9. R. B. Conley, Jr.
10. R. L. Davis
11. B. C. Duggins
12. J. L. Duncan
13. B. G. Eads
14. H. C. Ford
15-18. R. W. Hayes
19. G. D. Inman
20-27. J. M. Jansen, Jr.
28. J. J. Jernigan
29. W. E. Knoch
30. K. G. Leihkauff
31. A. M. Lokey
32. L. K. Lugten
33. D. R. Miller
34. C. A. Mossman
35. J. A. Murray, Jr.
36. A. B. Ralph
37. J. A. Russell, Jr.
38. K. J. Scott
39. P. Singleton
40. C. A. Smith
41. H. F. Smith, Jr.
42. J. B. Smith
43. W. W. Smith, Jr.
44. M. L. Stallo
45. B. A. Tye
46. R. S. Wiltshire
47. J. B. Ball (Advisor)
48. M. J. Kopp (Advisor)
49. P. F. McCrea (Advisor)
50. H. M. Paynter (Advisor)
51-52. Laboratory Records Department
53. Laboratory Records ORNL-RC
54. ORNL Patent Section
55. I&C Publications and Information Processing Center

56. Assistant Manager for Energy Research and Development, DOE-ORO, Oak Ridge, TN 37831.
57-161. Given distribution under Category UC-37, Instruments.