N92-12501

# Performance Analysis of LAN Bridges & Routers

Ankur R. Hajare, *MITRE Corporation*

## Abstract

Bridges and routers are used to interconnect Local Area Networks (LANs). The performance of these devices is important since they can become bottlenecks in large multi-segment networks. Performance metrics and a test methodology for bridges and routers have not been standardized. Performance data reported by vendors is not applicable to the actual scenarios encountered in an operational network. However, vendor-provided data can be used to calibrate models of bridges and routers that, along with other models, yield performance data for a network. Several tools are available for modelling bridges and routers, and Network II.5® was used for this study. The results of the analysis of some bridges and routers are presented in this paper.

# PERFORMANCE ANALYSIS OF LAN BRIDGES AND ROUTERS


## ABSTRACT

Bridges and routers are used to interconnect Local Area Networks (LANs). The performance of these devices is important since they can become bottlenecks in large multi-segment networks. Performance metrics and a test methodology for bridges and routers have not been standardized. Performance data reported by vendors is not applicable to the actual scenarios encountered in an operational network. However, vendor-provided data can be used to calibrate models of bridges and routers that, along with other models, yield performance data for a network. Several tools are available for modelling bridges and routers, and Network II.5® was used for this study. The results of the analysis of some bridges and routers are presented in this paper.


## INTRODUCTION

Bridges and routers are used to interconnect multiple segments of a Local Area Network (LAN). These devices reduce congestion on a LAN since they restrict traffic that is local to a segment while forwarding only those packets that are addressed to devices on other segments [Reddy, 1990]. As shown in figure 1, bridges operate at the Data Link layer, which is layer 2 of the 7-layer Open System Interconnection (OSI) model. A bridge examines the destination address field of all valid packets on a LAN segment and, using an address table for each segment, determines whether the packet needs to be forwarded [Backes, 1988]. A few years ago, bridges required explicit programming of their address tables before installation. Today almost all bridges are learning bridges, i.e. they generate their address table by themselves when installed in a network. Although a learning bridge is much easier to set up and manage, this convenience is achieved at the expense of performance.

As shown in figure 1, routers operate at the Network layer, which is layer 3 of the 7-layer OSI model. Thus, routers are specific to a protocol such as TCP/IP, DECnet or Novell IPX. Until about a year ago, routers could handle only a single protocol. However, vendors have recently introduced routers that can handle multiple protocols, even when they are intermixed. Routers examine the source and destination addresses and, in some cases, routing information within each packet. Since this information is regarded as data by the data link layer protocol, routers are insensitive to the layer 2 protocol that is being used. Routing imposes a larger computational burden on a device than bridging. Because of this, routers have performed slower than bridges. A performance ratio as high as 5:1 for bridging vs. routing has been reported [Spiner, 1990].

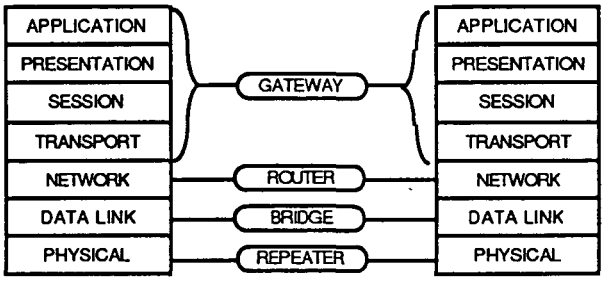


Figure 1: OSI Model Showing Repeaters, Bridges, Routers and Gateways.

Under certain circumstances, bridges and routers can become bottlenecks [Salwen et al, 1988]. Loss of packets by bridges and routers results in error conditions and re-transmission [Hordeski, 1987], which deteriorates end-user response times. Hence, it is important to measure and analyze bridge performance under various conditions that are encountered in an operational network [Rickert, 1990].

Most LAN performance studies focus on single segment performance [DuBois, 1988]. However, when end-to-end performance of a network is being assessed, bridge and router performance can be more important than the performance of the transmission medium [Boggs et al, 1988].

## RATIONALE FOR MODELLING

Vendors of bridges and routers provide performance specifications for their products. Since no standards presently exist for the specification of bridge and router performance [Jackson, 1989 and Salamone, 1990], different metrics are reported by different vendors. Information about the conditions under which the performance data was derived is generally not provided by vendors. Since the testing methodology is not standardized either, each vendor can create tests that demonstrate their own products to be superior [Bradner, 1991].

Although test results are available from several sources, the data provided is not directly applicable to a real situation. That is because the tests are performed under conditions that are not typical of what is encountered in actual network usage. Usually, tests are performed with all packets of one size that arrive at a steady rate. Consequently, the effect of differences in buffer sizes is not demonstrated. In contrast, LAN traffic in the real world is bursty and buffer size does affect performance. Furthermore, most reported measurements are performed for unidirectional forwarding of all packets in a single stream with no other traffic on the LAN. Such test results, though not directly usable, can be used to calibrate performance models of bridges and routers. The model can then predict performance for bursty, multiple data streams that contain a random mix of packets of various sizes.

Full scale testing of bridges and routers for a comprehensive set of scenarios is not practical because of the large amount of test equipment and effort that would be required [Bradner, 1991]. Therefore, modelling is a practical alternative to assessing end-to-end performance of a large multi-segment network.

The performance models described in this paper were part of an effort to build a discrete event simulation model of a campus wide multi-vendor, multi-protocol network planned at the NASA Johnson Space Center (JSC). As a part of the task of modelling this network, models of all the types of devices within the network were being considered. The data from some of them are presented here.

## MODELLING TOOLS

Performance models are either analytic models or simulation models. Several analytic models have been developed for single segment LANs [Stallings, 1987 and Boggs et al, 1988]. However, no adequate analytic models have been reported for inter-networking devices. Analytic models are based on assumptions that convert a real-world problem into one that is amenable to a closed-form solution. Simulation models, on the other hand, do not require such drastic or extensive assumptions.

Analytic models usually predict only steady-state conditions, whereas simulation models demonstrate the effects of transients and the effects of initialization. For example, a typical learning bridge rebuilds the address table every few minutes. Such transient conditions are best studied by means of a simulation model. Other transient conditions amenable to simulation modelling include broadcast packets creating a broadcast storm.

Simulation models can be developed using either a general purpose simulation language (such as GPSS or Simscript®) or a network modelling tool. General purpose simulation languages provide more flexibility and power but are harder to use. Network modelling tools enable quicker development of models but are relatively restricted in their capabilities. Examples of network modelling tools are Network II.5®, Lannet II.5®, Block Oriented Network Simulator™ (BONeS™), and LANSIM™. In addition to these commercially available tools, several large organizations, such as IBM and AT&T, have their own modelling tools for in-house use [Van Norman, 1988].

The tool used for this study was Network II.5®., which is marketed by CACI Products, Inc. of La Jolla, California. This tool is installed on an IBM compatible mainframe at JSC and is accessible by the user community via the Center Information Network (CIN). This study does not imply an endorsement of the tool by NASA or by MITRE.

Network II.5® builds a discrete event simulation model from a model definition consisting of basic entities that include processing elements, storage devices, transfer devices, and software modules. Each processing element has a set of instructions. Software modules, which consist of instructions, run on processing elements. These modules have fixed or probabilistic execution times. Processing elements can send messages via transfer devices to other processing elements or to storage devices. Messages queue at processing elements where they are processed by software modules. Also, software modules can queue for execution on processing elements. Network II.5® provides information on queue lengths and queueing delays, and it features scheduling mechanisms and priority disciplines. A random number generator and most of the commonly used statistical distributions are built into Network II.5®. Although Network II.5® is written in Simscript II.5®, no interface is provided to user-written Simscript II.5® code. A description of Network II.5® is provided by CACI [CACI, 1989].

Network II.5® contains built-in models for transfer devices that use collision, token ring, and other protocols. A specific LAN segment is, therefore, modelled by an appropriate selection of parameters. In addition to the built-in network protocols, Network II.5® provides the primitives necessary to model networking devices such as bridges, routers,

gateways, communications controllers, and front-end processors.

Network II.5® does not model at the physical layer. Thus, it does not model signal propagation along with phase shift, jitter, and error conditions. Network II.5® has a fixed sized collision window for each Ethernet® segment, whereas in reality it is a function of distance. Also, the inter-frame gap is fixed for a LAN. Thus, Network II.5® cannot handle variations in Network Interface Unit (NIU) speed that result in varying inter-frame gaps [Rickert, 1990].

## BRIDGE AND ROUTER ARCHITECTURE

Bridges and routers, typically, are microcomputer based and use a common chip, such as the Intel 80286® or the Motorola 68020®. They generally use a standard bus, such as VME® or Multibus®, which accommodates processor and memory modules, as well as the NIUs. Figure 2 illustrates the typical architecture used for bridges and routers. There are variations on this basic architecture, such as memory on the NIU board itself. Although an advantage in that the board provides additional memory, such an architecture can actually perform slower because the processor may be required to move data from the memory on one NIU to the memory on the other NIU.

A different type of router architecture that has been introduced recently is a dual-bus architecture, illustrated in figure 3. High-speed NIUs are interfaced to a high-speed bus, whereas slower NIUs are connected to a slower bus. Since simultaneous transfers can be performed on each bus, the performance threshold of the router is higher than a single bus architecture. A reason for retaining the slower bus (instead of using two high-speed buses) is to provide upward compatibility from older products that could only interface to the slower bus.

Vendors have recently introduced high-end products based on a distributed processing architecture, as illustrated in figure 4. The processor is usually the bottleneck in single processor designs, such as that of figure 2. Hence, performance can be improved either by a more powerful single processor or with multiple processors. Since the latter provides a higher performance threshold than the most powerful single microprocessor, vendors have recently come out with high-end routers based on distributed processing.

In the architecture of figure 4, the CPU performs control and monitoring functions. Although it may initiate transfers, the CPU does not participate in the actual data transfers between NIUs. Traffic between LANs that are connected to the same board in the router does not use the bus. Such multiple transfers can occur simultaneously without contending for resources, except for use of the CPU for initialization. Traffic between LANs that are connected to different boards does use the bus. Although the bus can interleave multiple transfers, there is contention for bus access, and this can limit throughput.

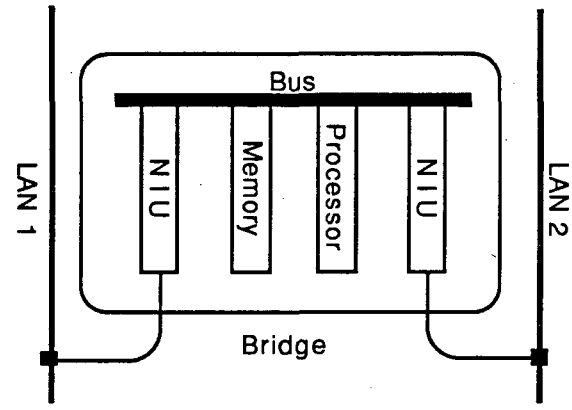


Figure 2: Typical Bridge/Router Architecture

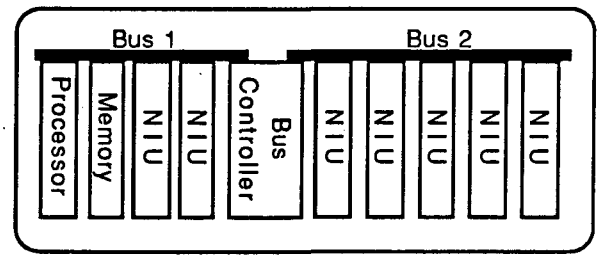


Figure 3: Dual-Bus Router Architecture

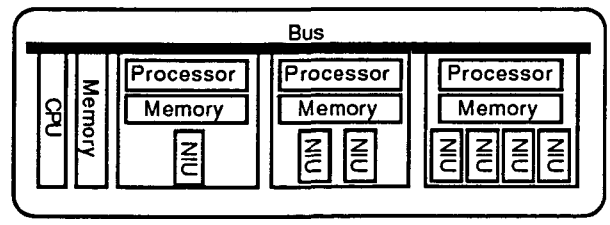


Figure 4: Distributed Processing Router Architecture

Although simple routers and bridges connect to just two LANs, the high-end products can connect several LANs. This has lead to their use as hubs [Korzeniowski, 1990], as shown in figure 5. Figure 6 shows an expanded view of a router configured to perform as a hub that interconnects one FDDI, one token ring, and four Ethernet LANs. In such a configuration, the bus of the router serves as the backbone. With a 32-bit bus, a transfer rate in excess of half a Gigabit/sec is claimed [Desmond, 1990].

## PERFORMANCE MODELS

Performance models of bridges and routers were developed using Network II.5®, based on vendor-provided information about the architecture and performance of each device. Given the architecture, its translation into Network II.5® terms was fairly straightforward in most cases. Buses were modelled as Network II.5® transfer devices, processors as Network II.5® processing elements, and NIUs were modelled as processing elements with buffer memory and I/O delays. Packet generation was by means of a Poisson process built into Network II.5®. The models were calibrated using reported performance data. Since several parameters were adjusted, many simulation runs were required for each model.

The data collected from the simulation runs included queue lengths, packet transfer times, and utilization of various resources such as processors, buses, and LANs. Due to the limited graphics capability and report generation capability of Network II.5®, it was sometimes necessary to use other software

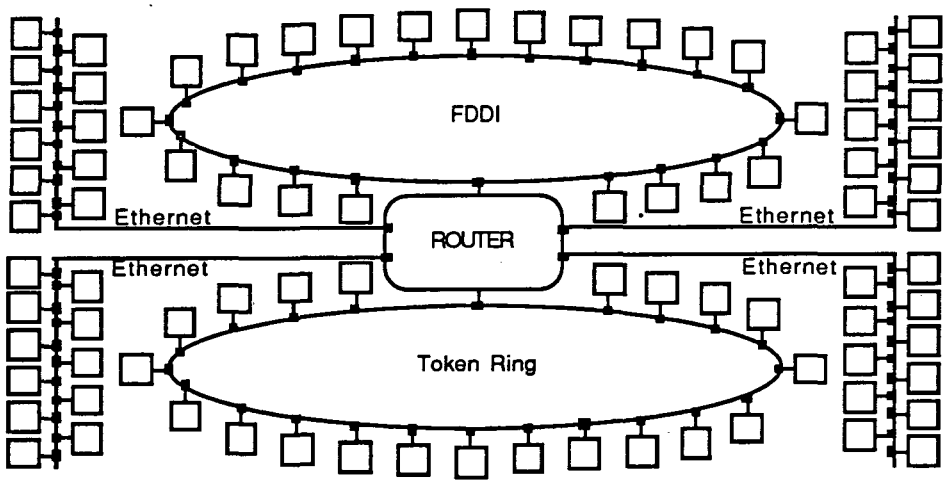


Figure 5: A Router as a Hub

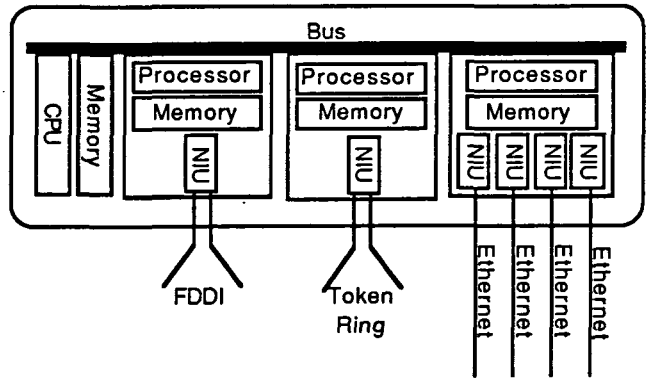


Figure 6: Configuration of a Router as a Hub

packages to analyze, format, and present the data generated by Network II.5®.

## RESULTS

The results of the performance analysis of some devices are presented here. The first of them is an Ethernet bridge. The processor in the bridge was a Motorola 68020® running at 20 MHz. The bridge used a Multibus® to connect the processor, memory, and two NIUs. It ran a Unix® kernel, optimized specifically for the device. The maximum unidirectional scan rate of the bridge was specified as 14K packets/sec, and the maximum bidirectional scan rate was listed as 22K packets/sec. The maximum forwarding rate was listed as 10K packets/sec. The packet delay, defined as the time from the end of packet reception to the start of packet transmission, was specified to be 150 µs. These performance specifications were used to calibrate the model. Bridge performance was studied for packet sizes ranging from the Ethernet minimum of 46 data bytes to the Ethernet maximum of 1500 data bytes. Several scenarios were investigated, and one of them is presented here.

Figures 7(a) and 7(b) illustrate the scenario where the bridge is forwarding packets in both directions. In this case both LANs had a random mix of packets, 50% of which had to be forwarded across the bridge. The maximum bidirectional forwarding rate that was achieved was 5800 packets/sec, in contrast to the vendor-rated 10,000 packets/sec. When packets arrived faster than 5800 packets/sec, some of them would be lost. For maximum-size packets, the bridge forwarded 1600 packets/sec. However, the amount of data forwarded by the bridge increased with packet size. This is illustrated in figure 7(b).

Figure 8 shows the performance of three bridges. Bridge A was based on a Motorola 68000® running at 12 MHz, and its transfer rate was specified as 7000 packets/sec. Bridge B is the one presented earlier in figures 7(a) and 7(b). Bridge C is a recently introduced high performance bridge with a multiprocessor architecture that contains a Motorola 68030® CPU. Bridges A and B differ noticeably only for small packets. However, bridge C can forward at a higher rate than the others for all packet sizes.

The performance of two routers is illustrated in figures 9(a) and (b). Both routers were single protocol devices that routed TCP/IP over Ethernet. Both utilized a single processor and were based on an architecture like that in figure 2. Although the

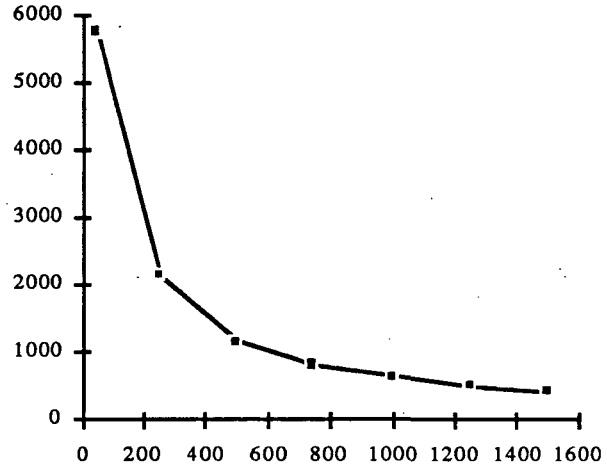


Figure 7(a): Bidirectional Bridge Throughput (packets/sec vs. packet size in bytes)

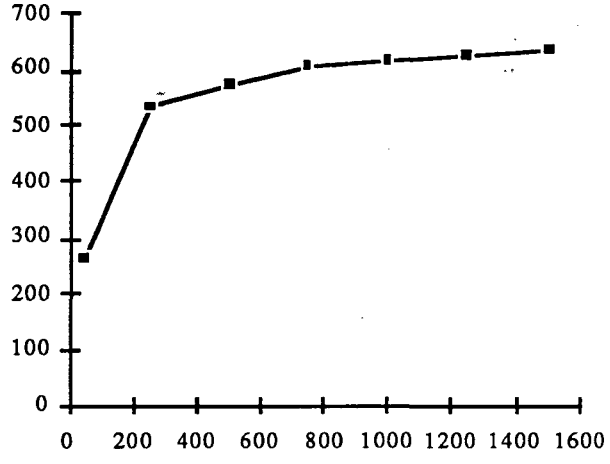


Figure 7(b): Bidirectional Bridge Throughput (Kbytes/sec vs. packet size in bytes)

routers could be configured with several Ethernet NIUs and were capable of routing multiple streams simultaneously, performance data was available only for routing a single stream. Figure 9(a) shows the unidirectional performance of the two routers in a scenario where all packets were forwarded and there was no other traffic on the two LANs connected to the router. As can be seen in the figure, the performance in terms of packets/sec decreased as packets size increased. However, as illustrated in figure 9(b), the amount of data forwarded by the router increased with packet size.

A router provides the capability to filter packets based on specified conditions, i.e. the router forwards only packets whose address information meets specified conditions. The conditions are based on a network management approach and are entered into a router when it is configured for operation. Checking filter conditions imposes an additional burden on the router and can affect its performance. This is illustrated in figure 10, which shows the performance of a router without filters, with one filter, and with ten filters. The router whose performance is shown in figure 10 is different from, and faster than, the ones whose performance is shown in figures 9(a) and 9(b).

Routers with a distributed processing architecture (as shown in figure 4) forward packets at different rates depending upon whether the forwarding is performed within a board or whether it is performed across boards. In the latter case, the data must be forwarded on the bus and, depending upon the router software, the process may impose a larger burden on the CPU. The performance of such a router is shown in figure 11. As can be seen in the figure, this router performs consistently better when forwarding packets within a board than for forwarding packets from one board to another.

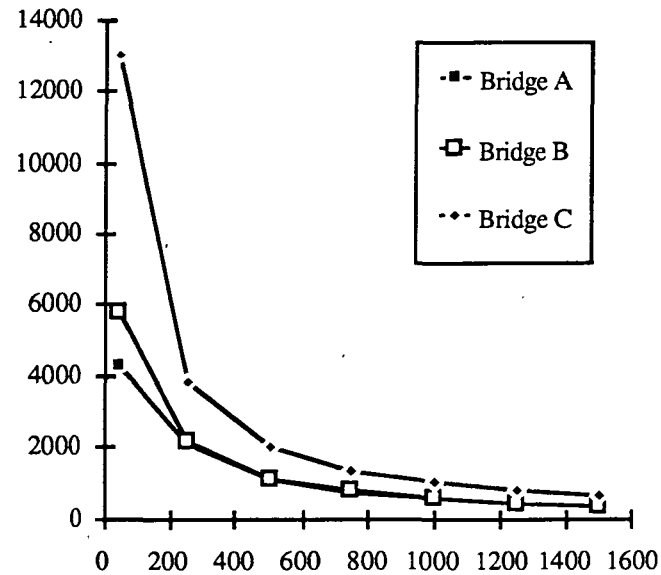


Figure 8: Bidirectional Throughput of Three Bridges (packets/sec vs. packet size in bytes)

## CONCLUDING REMARKS

The rationale for modelling bridges and routers has been presented in this paper. The tool used for the study has been described, along with the architectural considerations of bridges and routers that are pertinent to modelling. The results of the performance analysis of some bridges and routers have been presented. Performance data, such as that presented here, can be used in selecting bridges and routers. Models, like the ones described here, can be incorporated into an integrated network model that predicts various aspects of network performance for the wide range of conditions that are encountered in actual operation. The model can be used to assess the impact of changes in network configuration, including the selection and configuration of bridges and/or routers within a network.

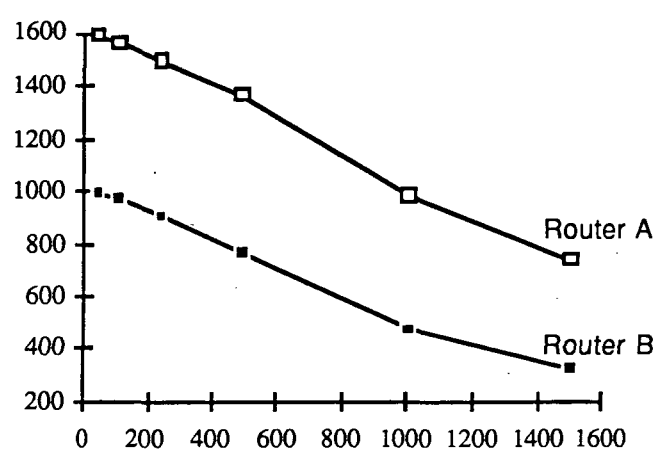


Figure 9(a): Router Performance (packets/sec vs. packet size in bytes)

## ACKNOWLEDGEMENT

This work was sponsored by NASA contract number NAS9-18057. Some of the performance data used in this study was provided by Scott O. Bradner of Harvard University.

## REFERENCES

Backes, F., "Transparent Bridges for Interconnection of IEEE 802 LANS", IEEE Network, Vol. 2, No. 1, ISSN 0890-8044, pp. 5-9, The Institute of Electrical and Electronics Engineers, New York, NY, 1988.

Boggs, D. R. et al, "Measured Capacity of an Ethernet: Myths and Reality", Proceedings of the Sigcomm '88 Symposium, Association for Computing Machinery, New York, NY, 1988.

Bradner, S. O., "Testing Multiprotocol Routers", Data Communications, Vol. 20, No. 2, pp. 71-86, McGraw-Hill, New York, NY, February 1991.

CACI, "Network II.5® User's Manual", CACI Products Company, La Jolla, CA, 1989.

Desmond, P., "New LAN Backbone Designs Simplify Net Management", Network World, Vol. 7, No. 44, International Data Group, Farmingham, MA, October 29, 1990.

Hordeski, M. F., "Microcomputer LANs: Network Design and Implementation", TAB Professional and Reference Books, Blue Ridge Summit, PA, 1987.

Jackson, K., "Router Standard Near?", Communications Week, November 27, 1989, CMP Publications, Inc., Manhasset, NY, 1989.

Korzeniowski, P., "The Fastest Backbone Net?", Communications Week, September 3, 1990, CMP Publications, Inc., Manhasset, NY, 1990.

Reddy, S., "Span the LAN", LAN Magazine, pp. 90-98, February 1990, Miller Freeman Publications, San Francisco, CA, 1990.

Rickert, J. R. Jr., "Evaluating MAC-Layer Bridges", Data Communications, pp. 117-122, May 1990, McGraw-Hill Inc., New York, NY, 1990.

Salamone, S., "The Need for LAN Benchmarks", Network World, Vol. 7, No. 37, Network World, Inc., Farmingham, MA, 1990.

Salwen, H. et al, "Examination of the Applicability of Router and Bridging Techniques", IEEE Network, Vol. 2, No. 1, ISSN 0890-8044, pp. 77-80, The Institute of Electrical and Electronics Engineers, New York, NY, 1988.


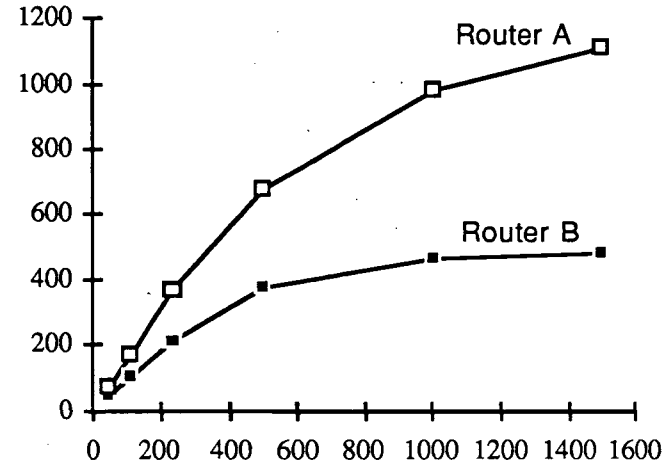


Figure 9(b): Router Throughput
(Kbytes/sec vs. packet size in bytes)

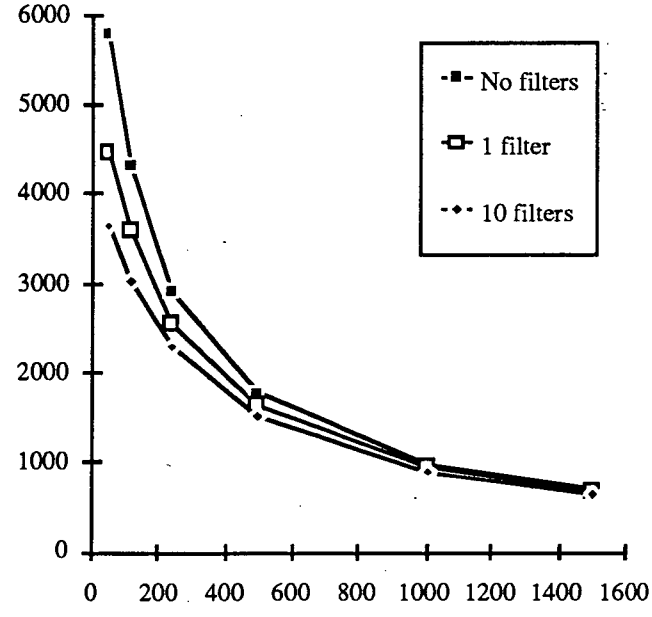


Figure 10: Router Performance with Filters
(packets/sec vs. packet size in bytes)


Spiner, D., "Internetwork Now!", LAN Technology, October 1990, M&T Publishing, Inc., Redwood City, CA, 1990.

Stallings, W., "Local Networks", 2nd edition, Macmillan Publishing Company, New York, NY, 1987.

Van Norman, H. J., "A User's Guide to Network Design Tools", Data Communications, April 1988.

LIST OF ABBREVIATIONS

| | |
|---|---|
| CIN | Center Information Network |
| CPU | Central Processing Unit |
| DEC | Digital Equipment Corporation |
| IPX | Internetwork Packet Exchange |
| NASA | National Aeronautics and Space Administration |
| MAC | Media Access Control |
| MHz | megahertz |
| µs | microseconds |
| NIU | Network Interface Unit |
| LAN | Local Area Network |
| OSI | Open Systems Interconnection |
| sec | second |
| TCP/IP | Transmission Control Protocol/Internet Protocol |

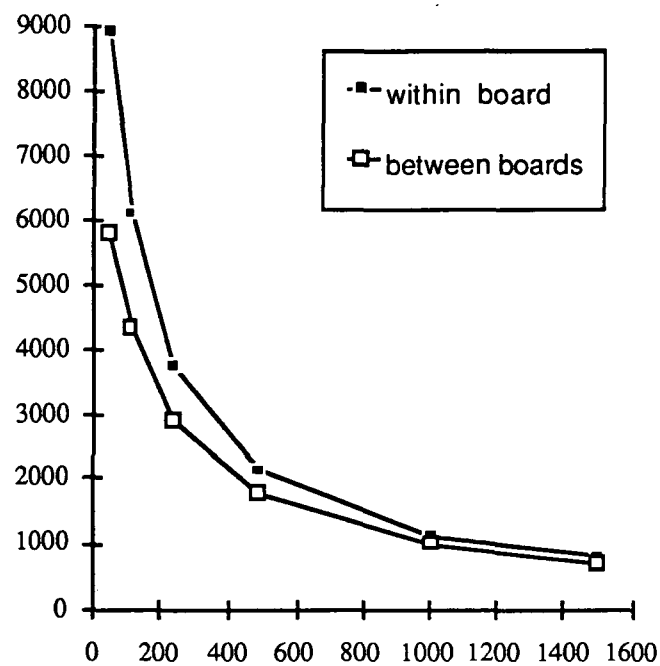


Figure 11: Router Performance
(packets/sec vs. packet size in bytes)