Proceedings of the 1991 Winter Simulation Conference
Barry L. Nelson, W. David Kelton, Gordon M. Clark (eds.)

INTERCONNECTED RING NETWORKS: A PERFORMANCE STUDY

Sub Ramakrishnan[1]
Pratap Thondapu

Bowling Green State University
Department of Computer Science
Bowling Green, Ohio 43403

ABSTRACT

This paper investigates the performance of multiple ring networks interconnected via bridges to a backbone ring. The networks employ IEEE 802.5 protocol. The model characterizes the behavior of the IEEE 802.2 protocol employed at the Logical Link Control Layer.

The primary performance measures are throughput and delay. The network carries only interring traffic. The model considers homogeneous traffic environments. It is shown that the protocol throughput is critically dependant on the window size and the buffer size at the intervening bridges.

1 INTRODUCTION

The technology of LANs is changing the way information is collected, processed and stored in organizations. LANs offer a favorable cost / performance ratio, graceful degradation, and support (emerging) distributed applications. As the importance of LANs grows, the requirements placed on these networks are exceeding the capacity of a single LAN. This has led, in turn, to a second generation of workstation connectivity, the linking of multiple LANs into an *internetwork*. This connectivity is desirable for a number of reasons, including: signal quality, availability, network diameter, addressing considerations, and security [Dixon and Pitt, 1988]. Interconnection may be achieved at any of the first three layers of ISO model using devices such as repeaters, bridges, routers, and gateways.

Numerous studies [Bux and Grillo, 1985; Gonslaves and Tobagi, 1988] have shown that the performance of interconnected networks largely depends on the internetwork configuration, medium access strategy and access protocol parameters (holding time and message priority for token ring, backoff interval and station distance for CSMA / CD networks). However, few studies have focussed on the link layer protocols for internetworks [Bux and Grillo, 1985; Biersack, 1988]. The IEEE standards 802.3, 802.4, and 802.5 govern respectively the media access schemes for bus, token bus, and token ring. The Logical Link Control (LLC) layer is described in IEEE 802.2.

The performance of (interconnected) LANs can be analyzed by using one of three contemporary techniques: simulation, analytical modeling, and actual measurements on operational networks. Although queueing techniques are powerful in analyzing simple networks, their use is somewhat limited in modeling complex systems because of the state space exploration. Performance of operational networks can be tuned / improved by actual measurement studies. It needs to be pointed out that carrying out actual measurements usually involves substantial time and effort. Simulation techniques, on the other hand, are suitable for modeling most systems. Of course, simulation usually demands quite a bit of computing time and some effort in the modeling and validation stages.

This paper concerns ring networks interconnected via bridges to a backbone ring. A simulation model is developed to characterize the performance of the 802.2 protocol at the LLC layer. It is assumed that the physical layer employes 802.5. As specified in the IEEE standards, the bridges operate below the Media Access Control layer and as such are transparent to protocols operating above or below the LLC layer.

The primary measures of performance are throughput and delay. It is shown that the protocol throughput is critically dependant on the window size and the buffer size at the intervening bridges.

The organization of the paper is as follows. In the next section, the elements of the internetwork and the LLC protocol as defined in 802.2 are described. We also provide some justifications for the simulation tool used. In Section 3, we state the assumptions, define performance measures of interest and the values

---

[1] This work was supported in part by NSF Grant CCR-9111292

for simulation parameters. Performance results and their interpretation are discussed in Section 4.

2 THE ENVIRONMENT

2.1 Network Description

An internetwork is specified by network topology, network traffic and network operation. In order to provide the motivation behind this development, we define these elements as used in this paper.

Hosts represent machines running user applications, and are the ultimate source / destination for messages. A network is the transmission facility used to transport messages. Bridges connect a network to the backbone. Several bridges are used to form an internetwork. Each of the networks and the backbone has a ring topology.

Network traffic is defined by the end points and whether the connection is message switched or circuit switched. Certain applications involve only limited amount of data exchange between source and destination hosts, where as other applications involve huge amounts of data transfer. Since LLC is a connection oriented protocol, in this study we consider only circuit switched networks. The protocol has three distinct stages: connection set up, data transfer and connection disconnect. Circuit-switching means that a path is setup between the source and destination, and all the data flows through this path. However, if a call is blocked because the path is unavailable, the call is retried after a retry interval. Once the call is set up data transfer state begins. During this stage the source host may send several messages to the destination hosts. These messages (are fragmented and) are sent over the network as one or more packets. The relevant parameters for packetizing can be specified in our model.

We now discuss the operational or dynamic parameters. The path for each circuit-switched call is defined by a specific routing strategy and includes a description of the network's routing algorithm. The routing algorithm determines how the internetwork traffic flows from the source host to the destination host. While general purpose networks may employ static or adaptive algorithms [Tanenbaum, 1981], the internetwork of interest in this paper employs a static routing algorithm. More precisely, the path packets travels through is statically defined at the beginning of the simulation run.

The inernetwork we model is given by Figure 1. At any instant only the node (host or bridge) holding the token is permitted to transmit packets onto the ring. The bridge receives and buffers packets (or frames) transmitted over the two rings it is connected to. Frame received on one ring is transmitted onto the other ring if the destination is on the path to the second ring.

2.2 Simulation Tool

The primary objective of the simulation model is to provide a flexible test-bed for the study of internetworks. Two types of simulation models can be envisioned: discrete-event and continous-event. The discrete-event model maps more closely to network simulation and requires less processing time. Therefore, we chose a discrete-event simulation tool. Several features are to be considered in selecting an appropriate simulation tool. (1) Instant prototyping - This will enable one to define the network and prototype it in a short time, cutting down on development time. One can concentrate more on the problem than on the nitty-gritty aspects of data communications. (2) Uniform representation - The model should provide minimum number of flexible *objects*, which, in our case, can be used to represent source, destination, and bridge. (3) Access protocol - It should be capable of modeling any access protocol, more specifically, the access schemes should adhere to the IEEE standard. Further, parameters like token slot time, token passing time should be configurable. (4) Traffic pattern - It should be possible to specify several source destination pairs, different types of traffic, and the statistical distributions for message sizes and call interarrival times.

General-purpose programming languages are flexible but require more time to develop a workable simulation model. Even the use of simulation languages such as Simula, SLAM requires a good amount of modeling skills due to the large number of parameters and options of a sophisticated protocol like 802.2. Therefore, it was decided to use a special purpose language / package that has built in support for the IEEE standards and various topologies. The candidate of choice was COMNET II.5 [CACI, 1990].

2.2.1 Comnet II.5

COMNET II.5 is a performance analysis tool for networks. Based on the description of a network and its routing algorithms, COMNET II.5 simulates the operation of the network and provides measures of network performance. It can be used to simulate networks using various switching disciplines. Use of COMNET II.5 usually calls for limited programming experience. Network descriptions are created with a convenient menu driven, window based user interface. COMNET 11.5 is structurally divided into two parts. *COMNETIN* is the menu-driven, screen editor that is used to create / modify a network description and consists of three major categories of data: network topology, network traffic, and network operation. *COMNET* is the network simulation program. The primary input to COMNET is the network description file created by COMNETIN.

Host attributes include number of packet-swiching processors, the packet-switching time per processor, and the total packet buffer space at the host. Subnet attributes include number of hosts connected to the channel group, the transmission speed(bit rate), and link-level (IEEE 802.2) protocol overhead. Moreover, a subnet can have access protocol attributes that determine how the connected nodes resolve contention for use of the same subnet; this attribute will also determine whether a subnet can represent a point-to-point channel, a polled multipoint line, a CSMA/CD or token passing LAN, or a random access radio channel.

2.3 IEEE 802.2

The 802.5 protocol employs multiple priority levels. To keep the protocol simple, we use only one priority level for all transmissions. The 802.2 protocol is a window based protocol. I frames (or packets) are numbered modulo-N. The sender may transmit up to W packets before waiting for an acknowledgement. Herein, W is the size of the window. Usually W = N-1. An acknowledgement numbered NR from destination indicates packets numbered NR-W, ... NR-1 have been received and that the destination expects the next insequence frame to be numbered NR. Frames received out of sequence are *not* buffered by destination. Therefore, it is necessary to (re-) transmit some frames even though the cause of the error is due to some lower numbered frame. It needs to be pointed out that the 802.2 protocol does not employ selective retransmission mechanisms. The sender employs a retransmission timer to figure out the status of unacknowleged frames. For full details, the reader is referred to [IEEE 802.5, 1985].

3 MODEL PARAMETERS

Fundamental to our development are the following practical assumptions. Each node experiences a non-zero processing time. The window size of the 802.2 protocol can be varied (though this is not the case with the standard). Further, there is a call set up time. (It needs to be pointed out that other researchers have not considered call set up time.) Recall that a message (received from the higher (user) layer for transmission over the internet to the destination node's user layer) may be broken up into multiple packets at the source. Acknowledgements are on a packet basis and not on message basis. However, even though packets are the transmitted frames, only messages have an interarrival distribution (see below). We believe that our modeling aspects are more characteristic of real networks. In the following, we specify values to some of these modeling parameters.

The network consists of three local rings and a backbone ring. All of them operate at 4 Mbps. We consider two types of configurations, viz: Type 1, and Type 2. In both types, the network traffic is only interring. The traffic pattern for the type 1 configuration is illustrated in Figures 2.

In type 2 configuration, each of the three satellite rings has 12 nodes. A1... A12 (Ring 1), B1...B12 (Ring 2), and C1... C12 (Ring 3). A1...A6 transmit data messages to B1...B6; B7...B12 transmit data messages to C1...C6; C7...C12 transmit data messages to A7...A12.

The packet processing times at host nodes and bridges are respectively 1 ms and 3 ms. The slot time (ms)) is set to 1000 ms. The IEEE 802.2 protocol is employed in the internetwork. Packet retransmission interval is assumed to be 250 ms. These parameters are consistent with the assumption made by other researchers (see [Bux and Grillo, 1985] for justification). The size of acknowledgement is assumed to be small (24 bytes). We do not model other types of overhead associated with data messages.

The message length (M) is geometrically distributed with a mean of 512 bytes. The maximum length of a packet transmitted on the ring is limited to 1024 bytes. Messages longer than 1024 bytes are fragmented into packets. The message interarrival time is exponentially distributed with a mean of λ. The value of λ is a function of the total offered load, L, and the number of sources, S. The relation is given by

$$L = (1/\lambda) * M * 8 \text{ bits/byte} * S.$$

For example, for type 1 configuration (S = 3), λ = .008 seconds (messages arrive every 8 milliseconds) for an offered load, L = 3 MBPS.

The system has been simulated for L = 1 MBPS, 3 MBPS, and 6 MBPS.

4 PERFORMANCE RESULTS

In the following we discuss the simulation results for the two traffic patterns: type 1 and type 2. We compute several quantities of interest. Throughput is defined as the number of bits of data messages transmitted per second (Kbps). The mean delay is the elapsed time from the origination of a message at the source and the receipt of its acknowledgment at the source. The utilization (for the backbone) is the percentage of time the backbone is busy carrying data messages. We also plot the buffer size B as a function of system parameters.

The input parameters that can be varied include bridge buffer size (B), window size (W), and offered load (L).

For offered loads 3 MBPS and under the type 1 configuration produced nearly 100 % throughput with an insignificant delay. This was so since the rings have sufficient bandwidth. It was observed that the mean buffer utilization at the bridges was below 10 %.

The variation of throughput against offered load is given in Figure 3(A). Note that the throughput equals the offered load for L <= 3 MBPS and drops as the offered load increases to 6 MBPS. Figure 3(B) illustrates the variation of throughput versus window size for a fixed offered load, 6 MBPS. The findings are consistent with our intuition that larger window size results in more traffic causing more congestion at the backbone.

The results for type 2 configuration are found in Figures 4 and 5. Figure 4(A) depicts the decrease in mean delay as the buffer size B is varied for a fixed window size, W = 5, and offered load, L = 3 MBPS. The corresponding throughput figures are given in Figure 4(B). It is interesting to note (see Figure 4(B)) that a small increase in buffer size yields good throughput figures. At B = 7 Kbytes the throughput is close to offered load (3 MBPS). Therefore, buffer size of 7K is adequate for this load. Similar comments can be made about Figure 4(A).

One would expect the buffer utilization to go down as the buffer size, B is increased. This is shown to be the case in Figure 5(A) for type 2 configuration. Further, as the buffer size is increased, the bridge discards fewer packets thereby resulting in *fewer* retransmissions at the link level protocol. In other words, an increase in buffer size results in a reduced load on the backbone, as shown in Figure 5(B).

The same behavior was observed by [Bux and Grillo, 1985]. However, in our simulation we also provided some insight into the buffer occupancy and the backbone utilization. Our simulation model is capable of handling other parameters such as message loss probability and effect of overhead on throughput.

5 CONCLUSION

In this paper we presented a simple simulation model of an interconnected ring network. Our findings suggest that an optimal buffer size exists for a given set of parameter values. Further increase in buffer size is not required. As observed by Bux and Grillo (1985), the window size can also be dynamically controlled to improve throughput.

We showed how COMNET II.5 [CACI, 1990] package can be used to set up the model. Our experience is that the user spends a substantial amount of time learning the intricacies of the package. Further, the package produces only certain types of outputs. It was difficult to control these parameters.

REFERENCES

Biersack, E. 1988. Techniken zum Zusammenschlub von... (in German). Ph.D thesis, Munich, Germany.

Bux, W and Grillo, D. 1985. Flow Control of Interconnected Token Rings. *IEEE Transactions on Communications:* 33(10), pp. 1058-1065.

CACI, 1990. COMNET II.5 Reference Manual, CACI Products Company, CA.

Dixon, R and Pitt, D. 1988. Addressing, Bridging, and Source Routing. *IEEE Network:* 2(1), pp. 25-32.

IEEE 802.2. 1986. Logical Link Control: ISO/DIS 8802/2. IEEE and John Wiley and Sons.

IEEE 802.5. 1986. Token Ring: ISO/DP 8802/5. IEEE and John Wiley and Sons.

Tanenbaum, A. 1981. Computer Networks. Prentice Hall.

Gonslaves, T and Tobagi, F. 1988. On the Performance of Station Locations and Access Protocol Parameters in Ethernet Networks. *IEEE Transactions on Communications:* 36(4), pp. 441-449.

AUTHOR BIOGRAPHIES

SUB RAMAKRISHNAN is an assistant professor of Computer Science at Bowling Green State University, Bowling Green, Ohio. From 1985 -'87 he held a visiting appointment with the University of Alberta, Edmonton, Alberta. He holds a Ph.D degree in Computer Science from Indian Institute of Technology, New Delhi. His research interests include network protocols and distributed computing.

PRATAP THONDAPU is a graduate student in the Department of Computer Science at Bowling Green State University, Bowling Green, Ohio. He holds a BS degree in Computer Engineering from the University of Pune, Pune.

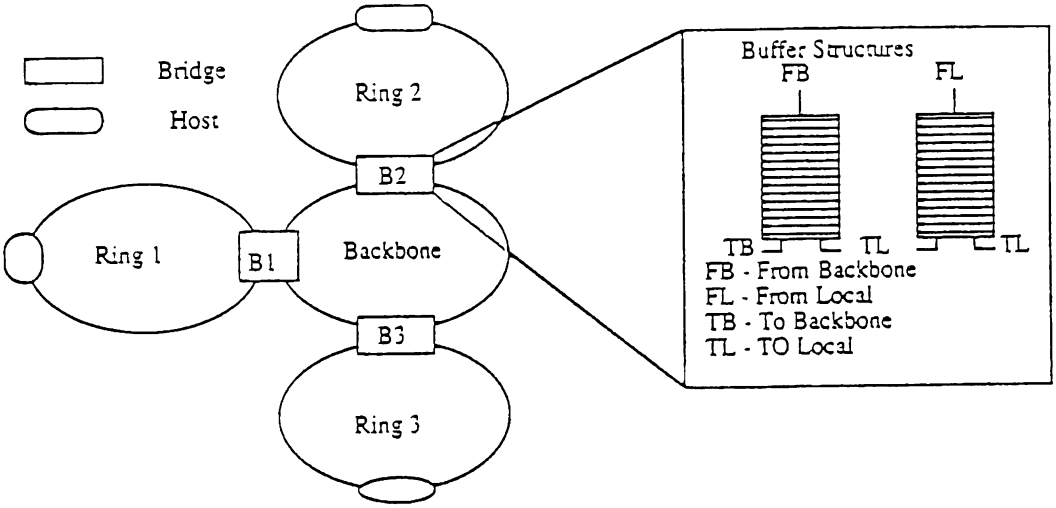

Figure 1 Elements of Internetwork

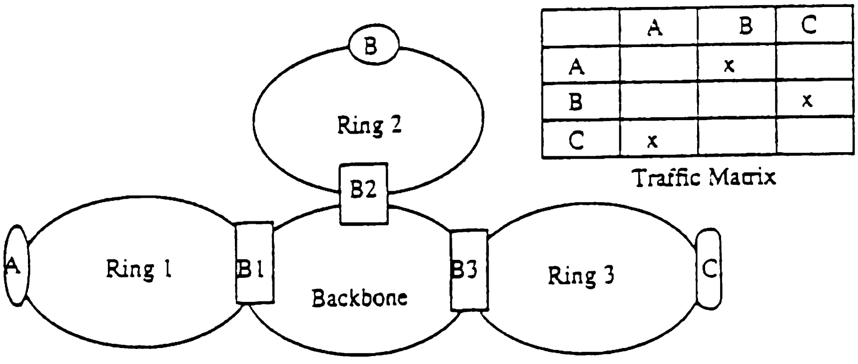

	A	B	C
A		x	
B			x
C	x		

Figure 2 Smal (Type 1) Configuration

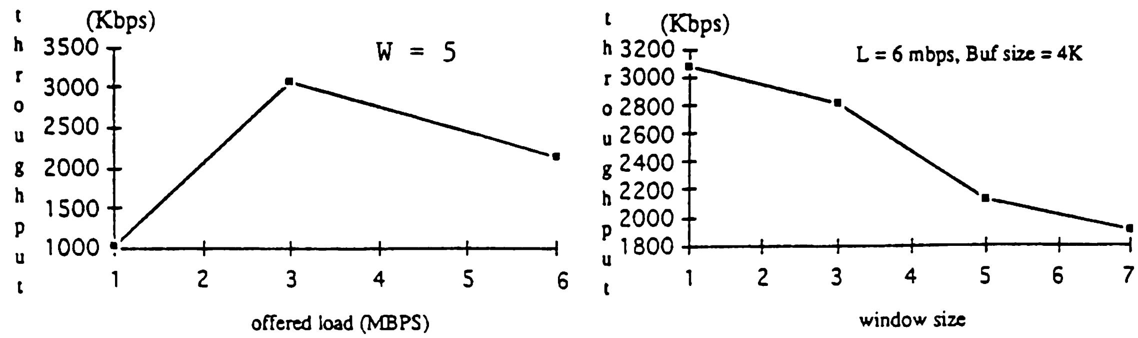

(A) Throughput vs offered load

(B) Throughput vs window size

Figure 3 Performance of Type 1 Configuration

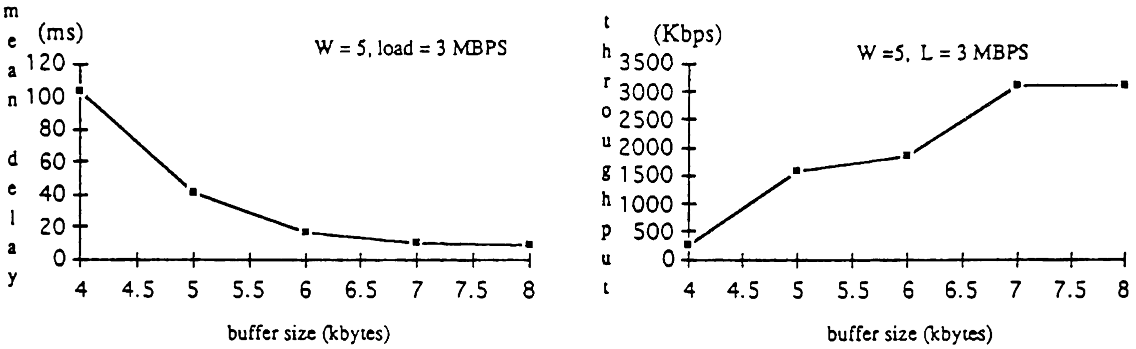

(A) Mean Delay vs buffer size

(B) Throughput vs buffer size

Figure 4 Performance of Type 2 Configuration

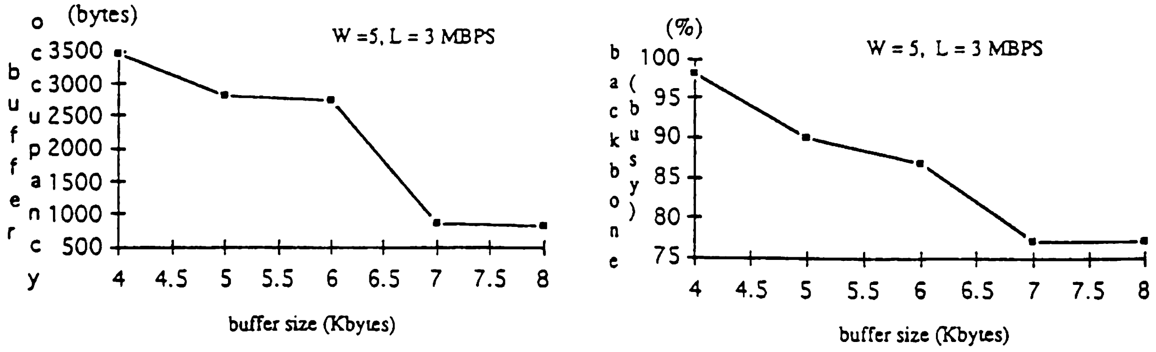

(A)

(B)

Figure 5 Effect of Buffer Size for Type 2 Configuration