Real-Time Performance Predictions: Transport Over LANs

K. Mills

M. Wheatley

S. Heatley

ABSTRACT

A five-layer simulation model of OSI protocols is applied to predict transport performance on a local area network (LAN). Emphasis is placed on time-critical applications typical of a small, flexible manufacturing system. The results predict that, with current technology, OSI protocols can provide one-way delays between 6 and 10 ms, and response times between 15 and 25 ms. The results also indicate that CSMA/CD is a reasonable access method for time-critical applications on small, factory LANS, if loads of less than 40% are anticipated. For loads between 40% and 70%, a token passing access method provides better performance for time-critical applications.

I. INTRODUCTION

To bridge the automation and information islands found in today's factories and offices, a significant portion of United States industry has selected the international standard open systems interconnection (OSI) protocols as described in the General Motor's manufacturing automation protocols (MAP) specification [MAP85] and the Boeing Computer Services' technical and office protocols (TOP) specification [TOP85]. The source for both

MAP and TOP is the OSI reference model [OSI82] and the series of protocol and service specifications that followed. Demonstrations at the 1984 National Computer Conference and the 1985 Automated Factory exhibition have shown that OSI protocols can provide effective interworking over multiple vendor networks and internetworks; however, questions remain concerning the efficiency of the protocols for planned applications and subnetwork technologies. What range of throughputs can be expected with present technology? What delays and response times can be achieved in factory and office applications? Are new protocol mechanisms required to meet performance needs? Must the seven layer architecture be modified for special time-critical applications? If so, what modifications are required? These questions and others are the subject of protocol performance research being conducted at the National Bureau of Standards (NBS).

This paper reports some early performance predictions concerning the OSI transport protocol, class 4, in use for several simple factory applications. The basis for these predictions is results obtained from a simulation model constructed to aid development of a plan for live performance experiments [HEA85A]. Simulation results are presented for periodic status reporting and request-response applications. For each application, transport performance is evaluated over both IEEE 802.3 and 802.4 networks [IEE84A, IEE84B, IEE85]. Some of the discussion relies on the key abbreviations shown in Table 1.

TABLE 1. KEY ABBREVIATIONS

TSDU	Transport Service Data Unit A user message.
IDU	Interface Data Unit TSDUs are divided into IDUs for passing between host and front-end.
TPDU	Transport Protocol Data Unit A message with a transport header.
DT	Data TPDU
AK	Acknowledgement TPDU

II. Status Reporting Performance Predictions

A common factory application is the periodic reporting of status between programmable controllers. In such applications, the primary performance concern is one-way delay. The experiment reported in the following paragraphs evaluates the expected one-way delay for a 20 millisecond periodic status reporting application between two transport stations. The local network background traffic varies from none to 70% in 5% increments and comprises 200 byte messages generated by a Poisson arrival process. One-way delay is measured as shown in Figure 1. A TSDU of 200 bytes is given to the sending transport at time Ts, encapsulated as a TPDU, and sent across the network to the receiving transport, arriving at the receiving user at time Te. The one-way delay (Te-Ts) is computed and accumulated for later calculation of the minimum, maximum, average, and standard deviation of the one-way delays measured on the transport connection. For each sample, 500 status report trials are used. The results are given in Figures 2, 3, and 4.

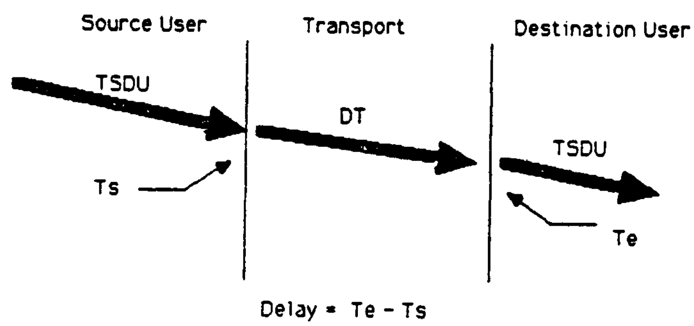

Figure 1. Measurement of One-Way Delay

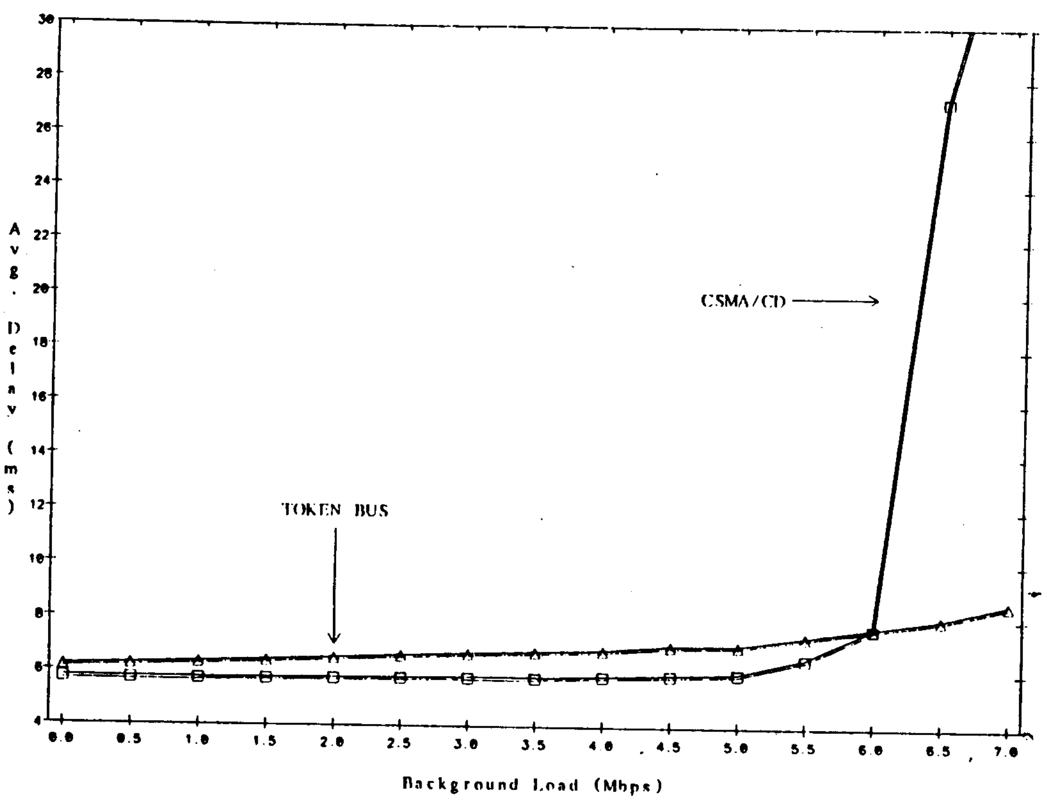

Figure 2. Transport User Average Delay for 20-Millisecond Status Reporting

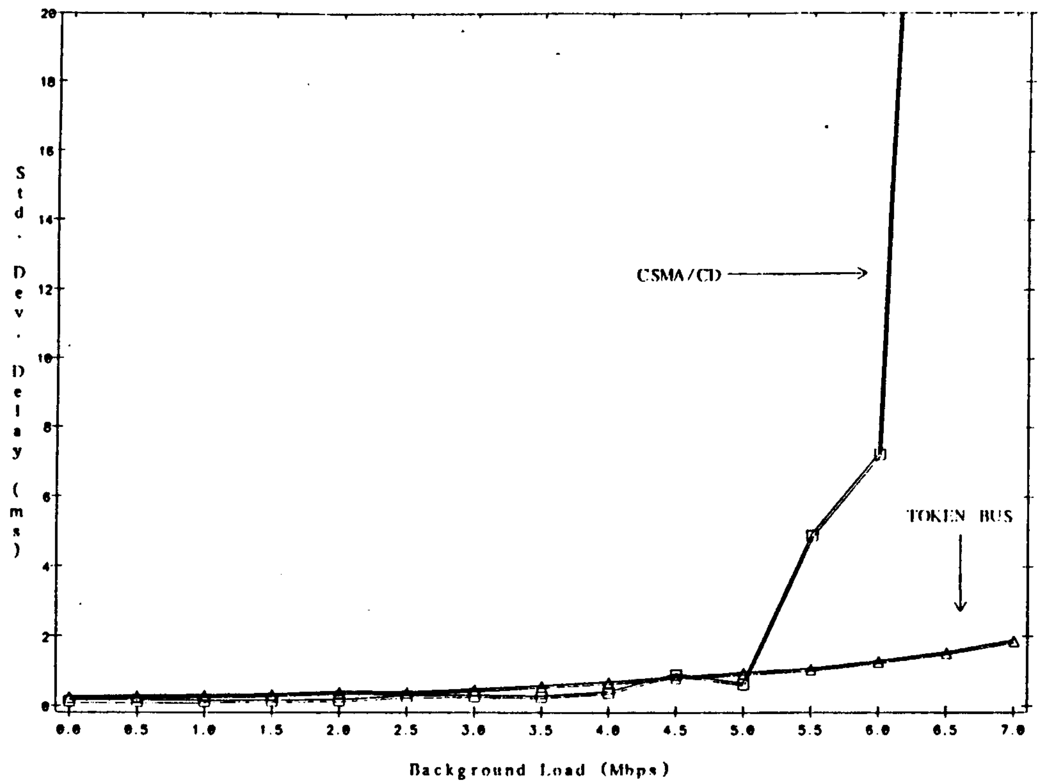

Figure 3. Transport User Standard Deviation in Delay for 20-Millisecond Status Reporting

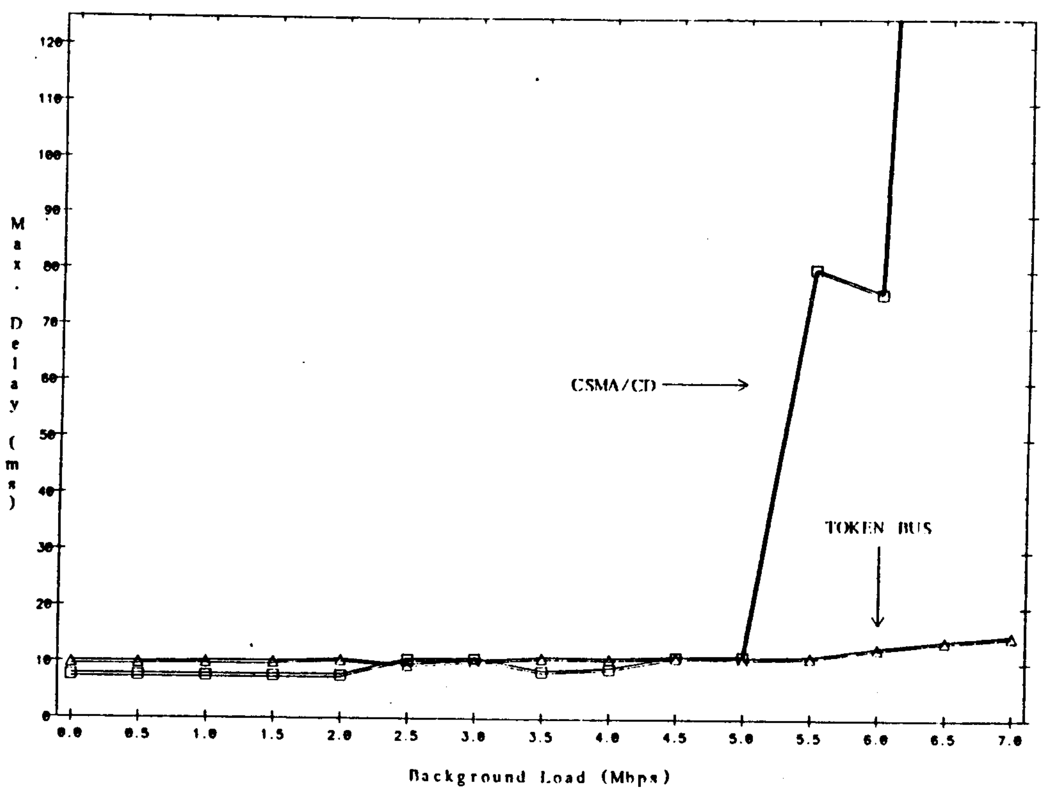

Figure 4. Transport User Maximum Delay for 20-Millisecond Status Reporting

The expected one-way delays for transport users over the CSMA/CD net are lower than the delays over the token bus until a cross over point is reached, at about 60% load, after which average one-way delays on the CSMA/CD net increase steeply. The user delays seen across the token bus, although somewhat higher for loads below 60%, increase in a more gradual manner. Therefore, if anticipated network loads are below 60%, the CSMA/CD local network provides better average one-way delay for the transport user.

If the application is sensitive to variance in one-way delays, Figure 3 shows that the advantage of CSMA/CD over token passing is reduced. The standard deviation in one-way delays is small when operating over both types of networks, up to a load of 50%, after which the deviation in one-way delays for the CSMA/CD network increases rapidly. Therefore, if loads of 50% or less are expected there is no clear advantage for either access method because the standard deviation in one-way delays is below 1 ms in all cases.

When the application must provide a guaranteed maximum one-way delay, Figure 4 illustrates that token passing has a potential advantage over CSMA/CD as the background load passes 50%; but, Figure 4 does not demonstrate the uncertainty associated with the CSMA/CD access method. An individual message can be substantially delayed due to repeated collisions. Therefore, one can expect significantly higher maximum delays with CSMA/CD at loads above 40%. Token passing provides a more controlled bound on the maximum one-way delay.

III. Request-Response Performance Predictions

Many factory applications entail requests for information and an associated response. The experiments detailed below examine the user reponse time obtained for such applications. In the first experiment, requests arrive at a constant rate without regard to previous responses. In the second experiment, a request is issued each time the response is received for a previous request. Figure 5 illustrates the method used to measure response time. The requesting user issues a 20 byte TSDU (i.e., the request) at time Treq, the TSDU is encapsulated as a TPDU and sent across the network to a responder. Upon receiving a request, the responding user submits a 200 byte TSDU (i.e., the response), the TSDU is formatted as a TPDU and sent across the network, arriving at the requester at time Tres. The reponse time is measured (Tres - Treq) and accumulated for later calculation of the mean and other distribution statistics. The experiments are conducted over both CSMA/CD and token bus LANs with background traffic varying from 0 to 70%. The composition of the background traffic is 200 byte messages generated by a Poisson arrival process.

A. Constant Rate

The constant rate experiment measures response times between a requesting and responding user over a single transport connection. The requester issues a new message every 20 ms. The responder issues a response immediately upon receipt of a request. The results obtained are given in in Figures 6, 7, and 8.

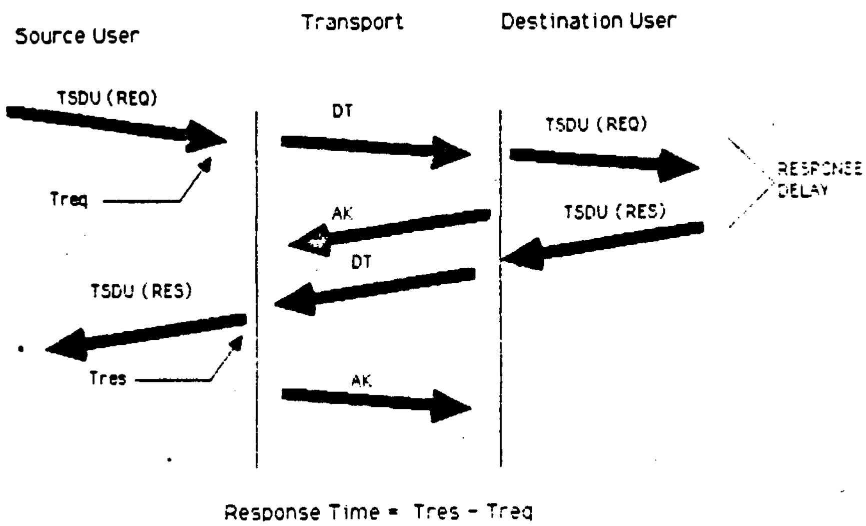

Figure 5. Measurement of Response Time

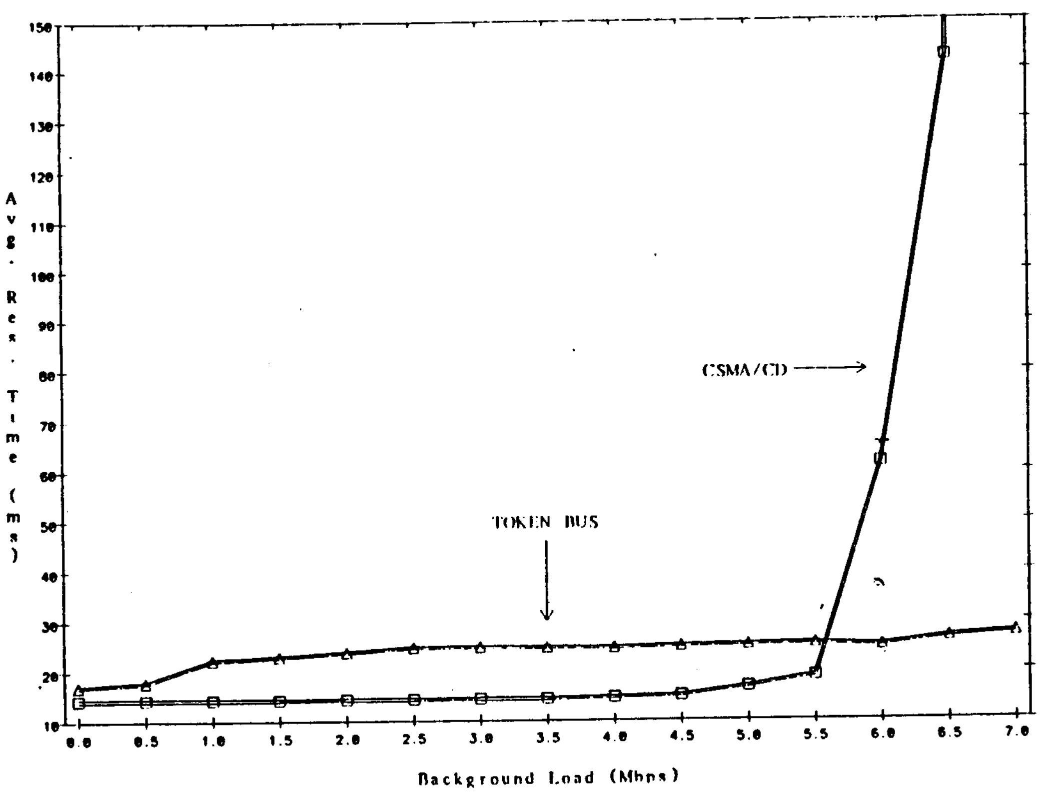

Figure 6. Transport User Average Response Time for Constant Rate Request-Response

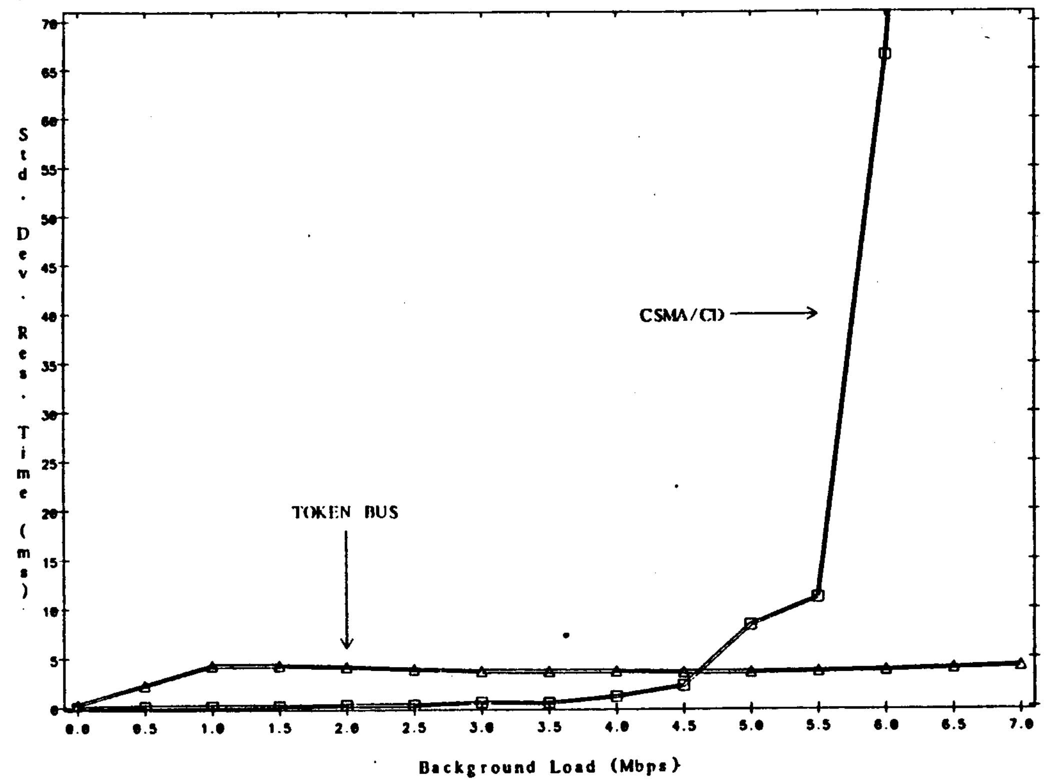

Figure 7. Transport User Standard Deviation of Response Time for Constant Rate Request-Response

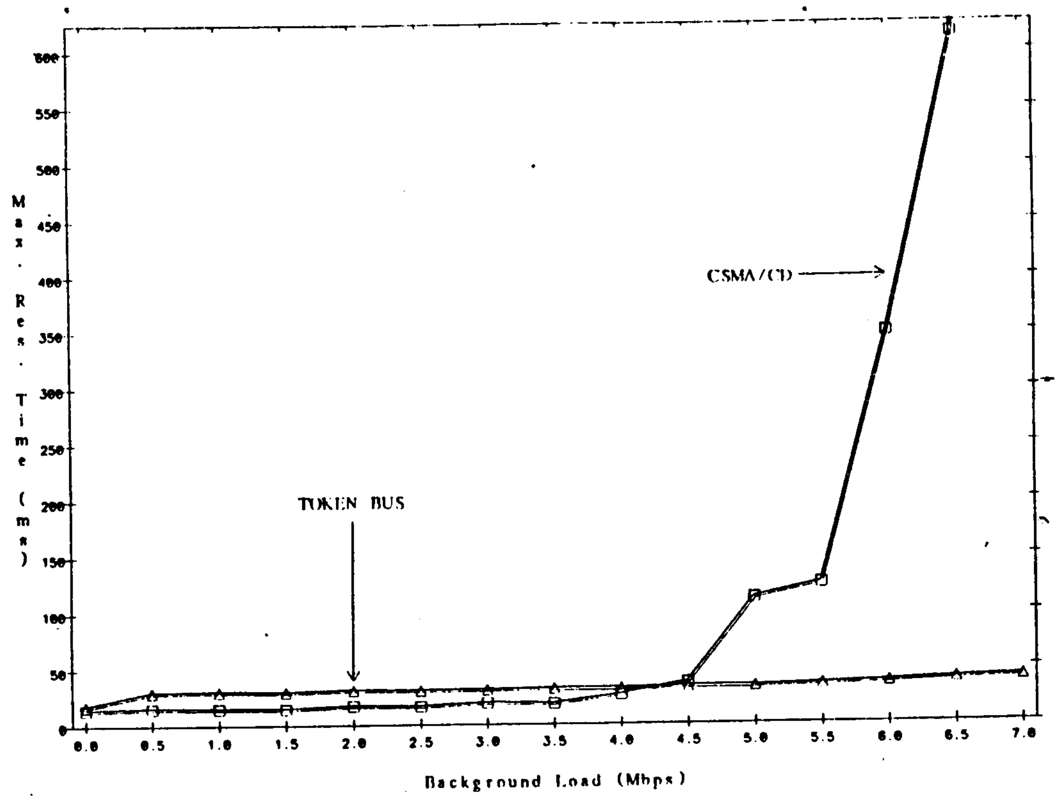

Figure 8. Transport User Maximum Response Time for Constant Rate Request-Response

For loads of 55% and below, the user sees superior average response times using CSMA/CD. If the application is more sensitive to variance in response time, CSMA/CD still provides superior performance but only for loads of 45% and below. Even for maximum response time the advantage of CSMA/CD is demonstrated for loads of 40% and below.

B. Maximum Rate

For the maximum rate experiment the requesting user submits requests as fast as responses are returned, thus, the message arrival rate adapts to the network load. The results are given in Figures 9, 10, and 11. At high loads, the average, standard deviation, and maximum response times are smaller than the same measures made with a constant arrival rate. The users only load the transport stations at the sustainable rate and no significant queuing occurs within the transport. In the previous experiment, as the background load increased, the users continued to load the transport station at a constant, nonsustainable rate and significant transport queuing occurred. Interestingly, at low loads, the average, maximum, and standard deviation of response times for the maximum rate experiment are sometimes higher than for the previous constant rate experiment. The explanation is simple. At low loads, where the average response time is below 18 ms, a dead time of 2 ms or more exists before the next request arrives. During this dead time, the two transport stations perform the processing associated with the AK TPDU for the response DT TPDU. When the maximum rate arrival scheme is used, no significant dead time exists between user receipt of a response and generation of the next request, thus the request must queue behind

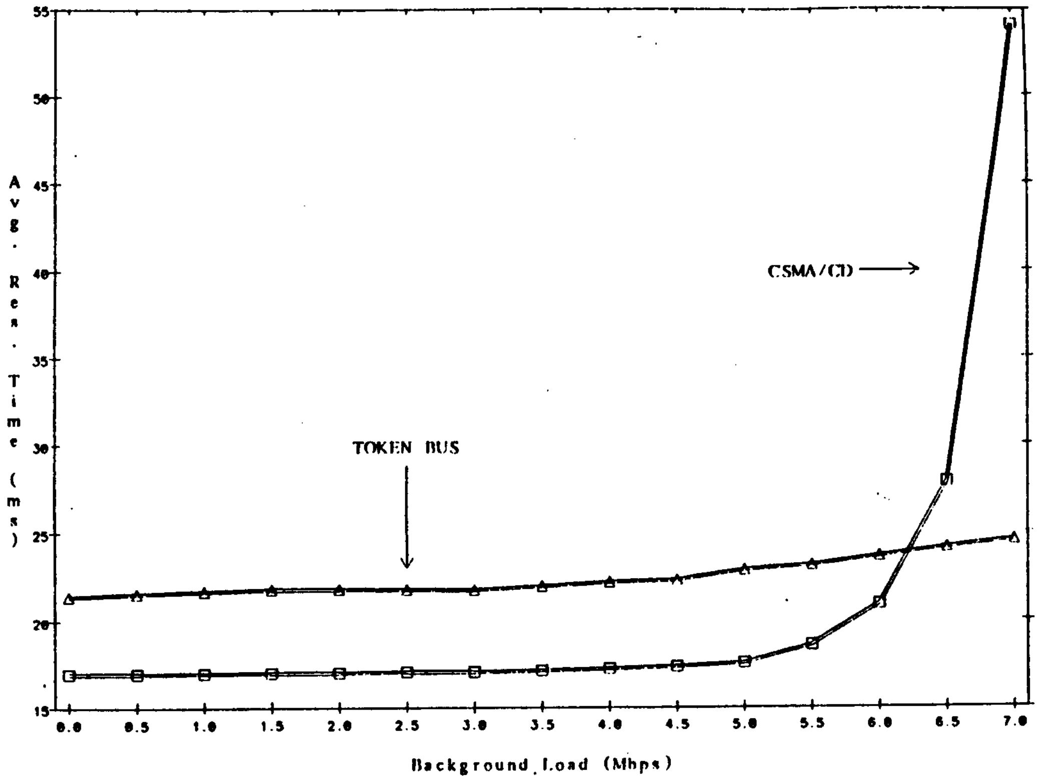

Figure 9. Transport User Average Response Time for Maximum Rate Request-Response

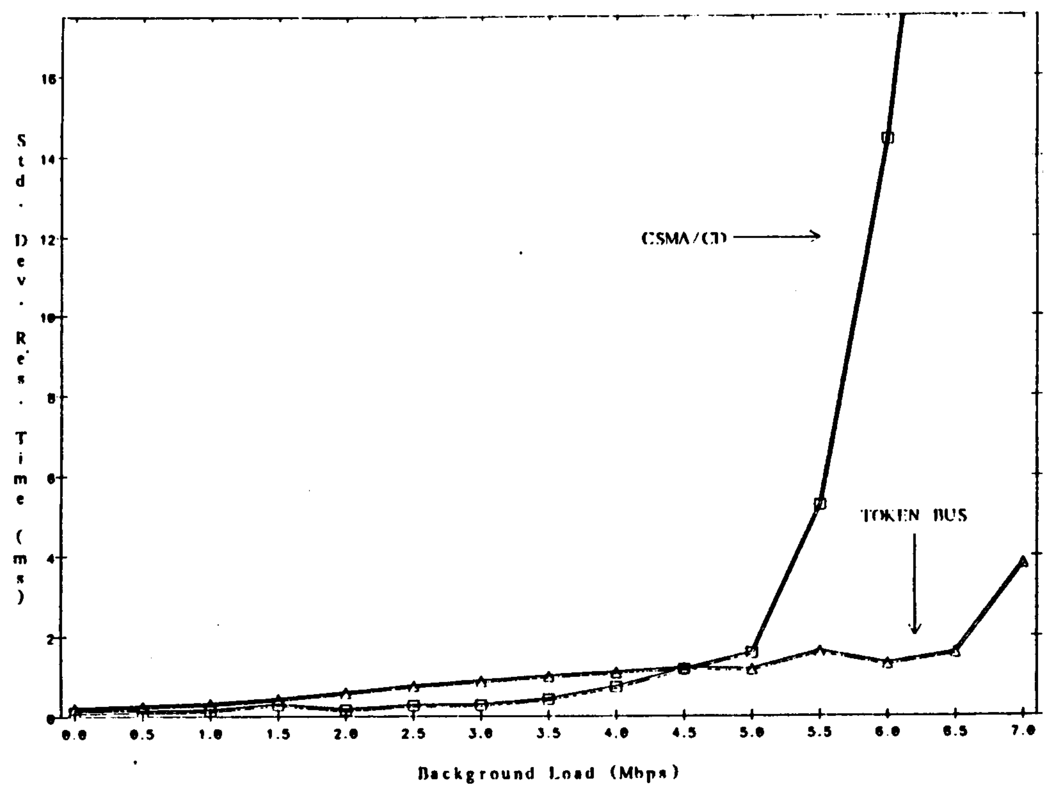

Figure 10. Transport User Standard Deviation of Response Time for Maximum Rate Request-Response

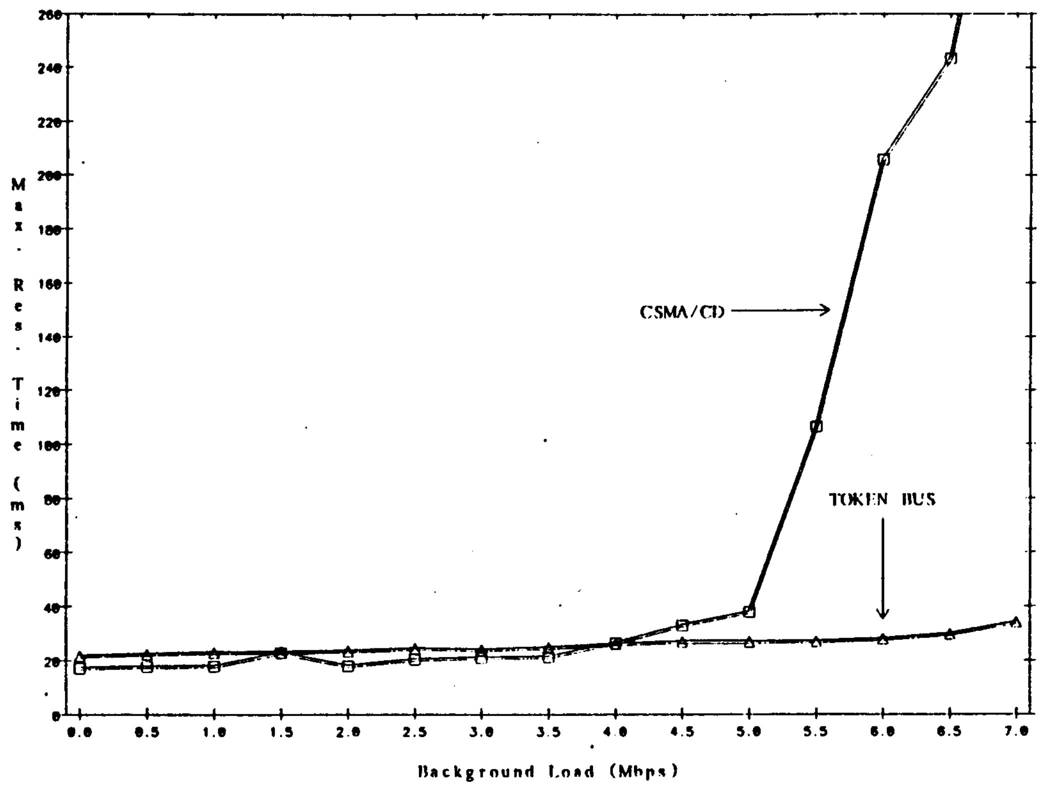

Figure 11. Transport User Maximum Response Time for Maximum Rate Request-Response

the AK TPDU for the previous response. This queuing occurs at both transports and a total of 2.7 ms is added to the one-way delay of the request message.

With maximum rate arrivals, the user sees better average response times over a CSMA/CD network for loads of 60% and below, but CSMA/CD provides no advantage over token passing when the measures of importance are standard deviation and maximum response times. For loads of 40% and below, CSMA/CD yields equal or better response time performance than token passing for maximum rate request-response applications.

IV. Conclusions

This paper describes simulation results from a detailed, five-layer simulation model of OSI protocols used to predict transport protocol performance. The findings, although not yet validated, suggest that OSI protocols, implemented with currently available technology and used on small, factory floor networks, can provide typical user one-way delays between 6 and 10 ms, and response times between 15 and 25 ms. These results, restricted to single connections with the traffic patterns as specified, are sufficient to encourage continued work to validate the simulation model through live experiments. The model is now being used at the NBS to plan experiments exploring multi-connection applications with varied user traffic.

With respect to choice of network access method, the results indiciate that for network loads of 40% and below, CSMA/CD provides performance equal to or better than token passing bus. This is shown for the average, standard deviation, and maximum response times and one-way delays. For small, flexible manufacturing networks, where loads below 40% can be expected the CSMA/CD access method can provide better performance for time-critical applications than token passing bus. When loads between 40% and 70% are anticipated token passing bus will yield superior time-critical performance.

REFERENCES

IEE84A IEEE 802.2 Standard for Logical Link Control, December 1984.

IEE84B IEEE 802.3 Standard for Carrier Sense Multiple Access with Collision Detection, December 1984.

IEE85 IEEE 802.4 Standard for Token-Passing Bus Access Method, February 1985.

HEA85A S. Heatley and R. Colella, Experiment Plan: ISO Transport Over An IEEE 802.3 Local Area Network, (under preparation).

MAP85 Maufacturing Automation Protocol, Version 2.1, General Motors, March 1985.

OSI82 Information Processing Systems - Open Systems Interconnection - Basic Reference Model, IS7498, International Organization for Standardization, 1982.

TOP85 Technical and Office Protocols Specification Document, Working Draft Version X.4.0, Boeing Computer Services, September 1985.