Performance Analysis of an ETHERNET Network Connected to a High-Speed LAN

J. Koutsonikos, V. Pallios, T. Antonakopoulos and V. Makios

First European Conference on Information Technology for Organisational Systems - EURINFO '88

ATHENS, MAY 1988

H.-J. Bullinger et al. (Editors)
Elsevier Science Publishers B.V. (North-Holland)

PERFORMANCE ANALYSIS OF AN ETHERNET NETWORK CONNECTED TO A HIGH-SPEED LAN

J.Koutsonikos D.Eng., V.Pallios M.Eng., T.Antonakopoulos D.Eng. and V.Makios Ph.D.

University of Patras, Laboratory of Electromagnetics

The Ethernet communication network is a broadcast, multiaccess system for local communications among locally distributed computing stations. This work presents the requirements and the results, through simulation analysis, for the interconnection of an Ethernet to a high-speed LAN. The asymmetric loading of the Ethernet, under various load conditions, which represents the connection to the network of a Gateway (GTW) is presented, and a comparison to the well analysed performance of the Ethernet in the literature is also attempted.

1. INTRODUCTION

The Ethernet network, introduced by R.M. Metcalf and D.R. Boggs [1], is a local communication network which uses a Carrier Sense Multiple Access with Collision Detection (CSMA/CD) access protocol on a tapped coaxial cable, to which all the communication devices are connected. It has a branching non-rooted tree topology and a 10 Mbs data rate. The 1-persistent CSMA/CD access protocol, used in the Ethernet, is devised in order to achieve acceptable throughput by never letting the channel go idle if some ready terminal is available [2].

The high-speed LANs under consideration have the following characteristics:

1) They support a wide variety of applications (e.g. office environment, universities, hospitals and industrial research laboratories),which implies that the system offers services both to stream type applications (such as telephony, videotelephony and still picture) and to bursty type applications (such as low, medium and high speed data).
2) The transmission rate is higher than 100 Mbs.

The simulation is based on the LION (Local Integrated Optical Network) network [3], which is a high speed LAN at 140 Mbits using a hybrid access protocol. The LION supports both circuit and packet traffic. The Gateway (GTW) between the Ethernet and the LION supports only packet traffic, directed from the LION to the Ethernet and vice versa. The GTW is a specially designed node of LION. This node of the LION is presented at the interface to the Ethernet as a "heavy" user, since the LION speed is much higher of the Ethernet speed and also the packet generation rate of the "heavy" user is greater than the packet generation rate of each of the other Ethernet users, due to the higher traffic supported by the LION.

The analysis of the performance of the Ethernet, under the previously mentioned conditions, is carried out through simulation techniques by producing the throughput-delay characteristics. The conditions under which the Ethernet is operating effectively are also discussed.

Section 2 presents the simulation model used for the analysis of the system. Section 3 includes the obtained numerical results of the simulation and a brief discussion. Finally, in section 4 the discussion of the Ethernet performance is addressed under the asymmetric loading.

2. THE SIMULATION MODEL

The simulation program, written in Pascal and developped on an IBM AT-3 personal computer, uses the technique of discrete event simulation [4,5].

The traffic source used, consists of a finite number of N connected users, communicating under the Ethernet, who collectively form an independent Poisson source with an aggregate mean packet generation rate of λ packets per second. One of these users will be the "heavy" user, with a higher packet generation rate, representing the Gateway between LION and Ethernet. The traffic source is distributed to the users in such a way, that the "heavy" user arrival rate is higher than the arrival rate of any other user by a factor of μ.

In the simulation model, each user has a buffer that can hold at most 3 packets, except the "heavy" user whose buffer can hold 10 packets. When a packet is generated by a user, it enters the user queue, under the condition the buffer is not full, and it is retained there until it is successfully transmitted. If the user´s buffer is full the packet is rejected.

A differentiation is introduced in the simulation model, between data arrived to a node and packets that this node transmits to the Ethernet. The former is refered as message coming to the node, which is a random variable exponentially distributed, with a mean length of l bytes per message. The procedure of transforming messages into proper packets for transmission includes segmentation of the message and addressing. Segmentation is taking place when the message length is greater than the maximum packet size (including the required control information) permitted for an Ethernet packet, which is 1518 bytes. The addressing includes the addition of a source address field for the sending node, a destination address field for the user for which the packet is directed to, a type field, and a frame check sequence for packet error detection. These fields have a total length of 48 bytes [6].

In the simulation program, each node is represented by a record, whose fields include the node number, the time at which this node will attempt to transmit the first packet of its buffer, two pointers (described in detail later), and the node statistics including the packets already tansmitted, the delay of these packets, the lost packets due to inadequate buffer space and the packets rejected due to 16 consequently failed attempts of each packet to be transmitted. The N records, one for each node, are linked together with the pointers mentioned before, in a manner that their transmission times have an increasing order. The result is called a linked ring, for which the head component is, at any time instant, the node who will first attempt to transmit a packet.

With the carrier detection feature, a given station will attempt to transmit when only it senses the channel idle. A station experiences a collision only during the initial part of its transmission attempt, the so called "collision window", thus the collision occurs before the transmitted signal has the time to propagate to all other nodes of the Ethernet channel. The time required for a node to acquire the channel is thus based on the round-trip propagation time of the physical channel.

The simulation program examines the transmission times of the two first records in the linked ring; if their difference is greater than the round-trip propagation time, the first user will successfully transmit his packet, otherwise a collision will occur. In the former case the statistics of the user who has been successfully transmitting are calculated, then a new transmission time for his record is scheduled, taking into account the Poisson arrival process and his buffer state and this record is placed in a new position in the ring. In the latter case, the new transmission times are scheduled for the users who participated in the collision, according to the retransmission algorithm described in the Ethernet version 2.0 [6], as the binary exponential backoff. Then the records of the users who participated in the collision are placed in new positions in the ring, in accordance to their new transmission times. The users for which the scheduled transmission times are contained inside the transmission epoch of the packet under transmission, are these that currently sense the channel busy; for these users, the transmission time is rescheduled at the end of the current transmission and their records are placed in the proper positions in the ring.

The parameters of the network, which are the end-to-end propagation time, the jam size, the interframe spacing, the backoff limit, the attempt limit, the slot time and the collision window, are described in the Ethernet, version 2.0.

3. SIMULATION RESULTS

The simulation results are presented in this section, based on the transmission of, at least, 200 packets from each Ethernet user, for reliabili-

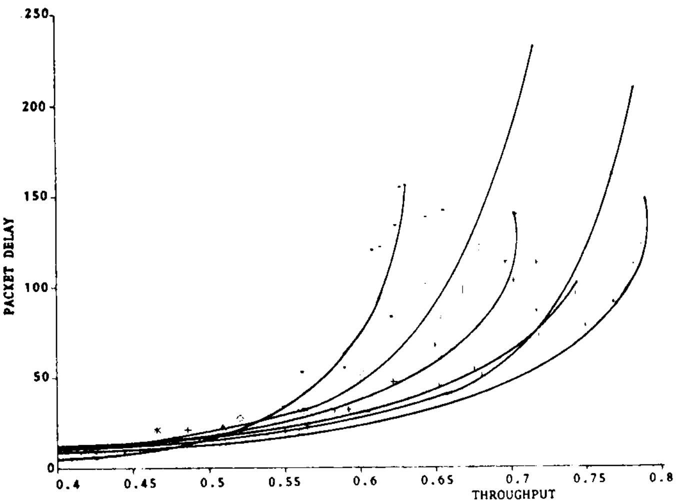

Fig.1: Throughput-Delay tradeoffs of the network for 10 users.

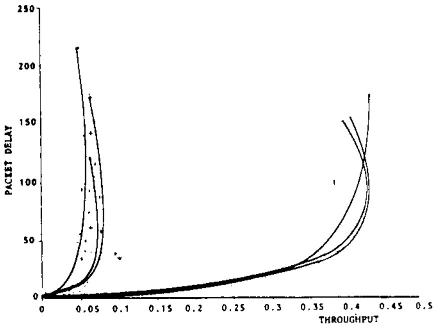

Fig.2: Performance of the "heavy" user,N=10.

ty purposes. The curves that will be exhibited and analysed, display the performance of Ethernet in terms of throughput-delay characteristics, under various parameter values.

The variable parameters are:

- N, the number of nodes connected to the Ethernet, taking the values 10 and 50.
- μ, the factor indicating how many times the mean arrival rate of the "heavy" user is greater than the mean arrival rate of the other users. The repertory of the μ values is 1,2,5 and 10; the value μ=1 represents the symmetric case, when the load is distributed fairly to all users.
- l, the mean message length, in bytes, taking the values 500,1000 and 2000.

At the following, the delay values are considered to be normalized; they represent the mean delay of each packet, in packet times. The mean message generation rate λ (which is also normalized in messages per mean message time) is less than 1, in all cases.

The examined cases are:

1) The performance of the network; in this case the throughput-delay characteristics of the network are presented, including all the Ethernet users.
2) The performance of the "heavy" user only, under various values of N, μ and l.
3) The performance of the mean "light" user. The term "light user" is used to distinguish it from the "heavy user", who has much higher λ. In an asymmetrically loaded network there are N-1 "light" users and one "heavy" user. The term "mean light user" is refered to all "light" users, normalized to one, e.g. mean "light" user´s throughput equals all "light" users´ throughput divided by N-1.

The results are divided in two parts; in the first part a comparison between the symmetric (μ=1) and the most asymmetric (μ=10) cases is presented, while in the second part a comparison between cases for all values of μ is given.

A) Comparison Between The Symmetric And The Most Asymmetric Cases.

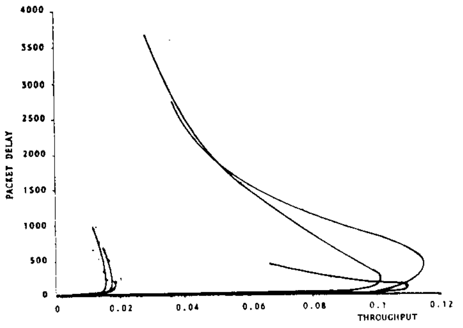

Fig.3: Performance of the "heavy" user, N=50.

Where the symbols in Figs. 1-5 are:

• is for μ=1 and l=500,
+ is for μ=1 and l=1000,
* is for μ=1 and l=2000,
o is for μ=10 and l=500,
x is for μ=10 and l=1000,
- is for μ=10 and l=2000.

The following subcases are distinguished:

1) For the performance of the network.
   In Fig.1 the throughput-delay characteristics of the network for N=10 is depicted. In this Figure, μ takes the values 1 and 10 and l takes the values 500,1000 and 2000. It is observed:
   - Throughput is higher and delay is lower in the asymmetric case, when 10 nodes are connected to the Ethernet. With N=50, the differences between symmetric and asymmetric loading are almost invisible.
   - There is an improvement in the performance of the Ethernet (in terms of greater achieved throughput and lower delay), when the mean message length increases. This difference is more sensible when 50 users are connected to the Ethernet.
   - Delay is much higher when N=50 than the case where N=10.
2) For the performance of the "heavy" user.
   In Figures 2 and 3 the throughput-delay curves of the "heavy" user are depicted, for N=10 and N=50, respectively. It is observed:
   - The "heavy" user throughput is much higher

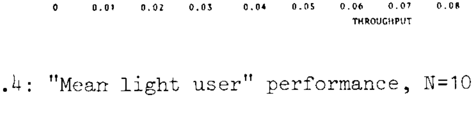

Fig.4: "Mean light user" performance, N=10.

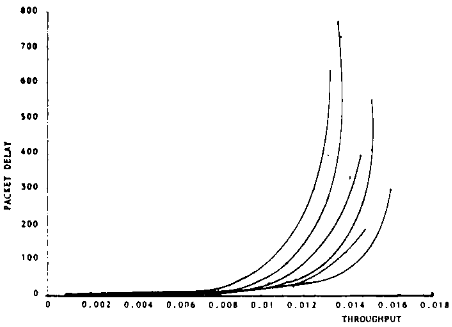

Fig.5: "Mean light user" performance, N=50.

in the asymmetric loading.

- The "heavy" user delay is extremely high when 50 users are connected to the Ethernet.
- In a system of 50 interconnected users, with λ (the total mean message generation rate) increasing up to 0.7, the "heavy" user throughput is also increasing up to its maximum value 0.12, in the asymmetric loading; for higher values of λ, the throughput decreases approaching the value of 0.02, which is the maximum achieved throughput of any user in the case that the load is symmetrically distributed to the users. This happens because, for high λ, any user, "heavy" or "light", has at any time instant, packets to transmit in his buffer and all users access the channel in the same way, resulting to an almost similar throughput for all users. A similar behaviour for the "heavy" user throughput is not observed in the case of 10 interconnected nodes, because λ is not so high in order to be able to reach such a situation.
- The delay is increasing when the mean message length increases.

3) For the performance of the mean "light" user. In Figures 4 and 5 the throughput-delay curves of the mean "light" user are depicted, for N=10 and N=50, respectively. It is observed:
   - The mean "light" user throughput is higher for the symmetrically loaded network. The difference is obviously larger in the case of 10 users.
   - The mean "light" user delay is lower in the asymmetric loading for 10 connected users, but the contrary happens in the case that N=50.
   - The mean "light" user performance is greater (in terms of greater throughput and lower delay) as the mean message length increases.
   - For 50 users connected under the Ethernet, the mean message length affects the mean "light" user more than μ does: the contrary happens for 10 Ethernet nodes.

B) Comparison Between Various Asymmetric Situations.

Where the symbols in Figs. 6,7 are:
• is for $\mu=1$,
+ is for $\mu=2$,
* is for $\mu=5$,
o is for $\mu=10$.

In this part, the differences in performance between various asymmetric cases (with μ taking the values 1,2,5 and 10) is presented. In Fig.7, the performance of the "heavy" user is predicted for N=50 and l=1000 bytes: in Fig.8, the performance of the "light" user is predicted for N=10 and l=2000 bytes. The observations presented following are produced not only considering the two cases previously mentioned, but also the cases with all possible combinations for N equals 10 and 50 and l equals 500, 1000 and 2000.

The following subcases are observed:

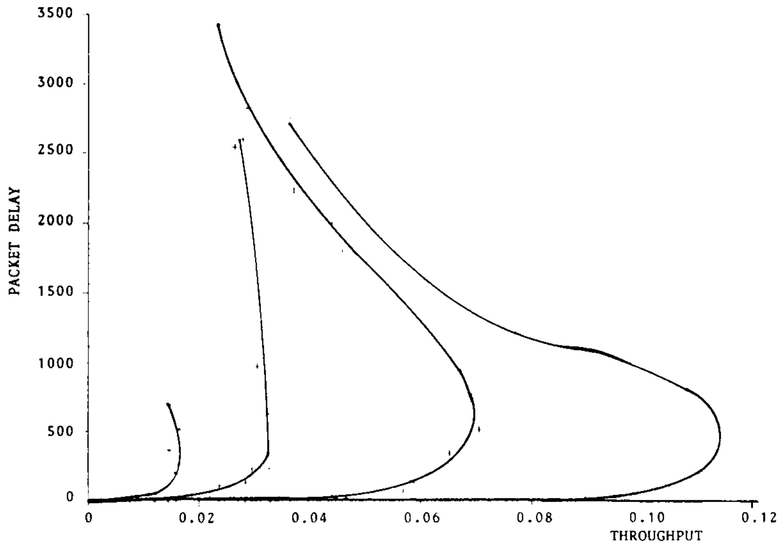

Fig.6: Throughput-Delay tradeoffs of the "heavy" user for N=50 and l=1000.

1) For the performance of the network, it is observed that:
   - The difference in throughput-delay curves between various values of µ are very small; specifically, for 50 connected users, the curves almost coincide.
   - A greater throughput with lower delay is achieved when the mean message length increases.
2) For the performance of the "heavy" user, it is observed that:
   - The "heavy" user throughput is clearly increasing with µ, for every N and l.
   - The "heavy" user delay is increasing with µ for N=50, but it is decreasing with µ for N=10; this happens because of the greater competition for the channel capture in the case of 50 interconnected users.
3) For the performance of the mean "light" user, it is observed that:
   - The throughput of the mean "light" user is decreasing when µ increases. The difference regarding throughput in various cases of µ is greater in the case that N=10. When N=50 the curves are very close, because the difference introduced by the "heavy" user to the throughput-delay characteristics is shared to a higher number of users.

Finally, considering the throughput-delay curve divided in two parts, named stable (where delay values are less than 50 packet times) and unstable (where delay values exceed 50 packet times), one could observe the following:

a) The delay is smaller, thus better, for smaller messages. This happens for any value of the variable parameters. The difference between the delay values for l equals 500,1000 and 2000 bytes is in the order of 1.5 packet times.
b) The factor µ is not affecting at all the throughput-delay characteristics of the network, but it is strongly affecting the performance of the "heavy" user. Regarding the mean "light" user, the affection of µ is sensible only in the case of 10 interconnected users.

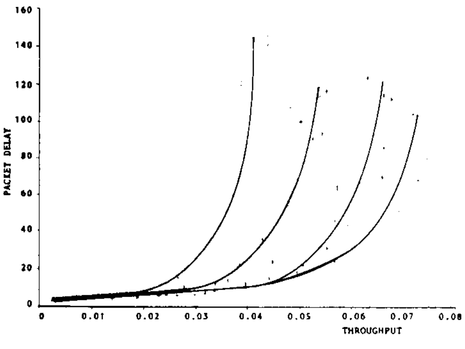

Fig.7: Throughput-delay tradeoffs for the "mean lightuser" N=10, l=2000.

4. CONCLUSION

A simulation analysis of the throughput-delay characteristics in an Ethernet network has been performed, under asymmetrically distributed load conditions. The introduction of the concept of the Gateway between the Ethernet and a high-speed Local Area Network gives rise to these asymmetric conditions.

The Gateway, represented by the "heavy" user in the simulation program, affects the Ethernet performance in several ways, depending on the system characteristics; the influence of the Gateway to the Ethernet network is decreasing as the number of nodes connected to the network increases.

The affection of the Gateway to the Ethernet is concentrated in the fact that the GTW increases the delay of the network packets. Regarding the stability of the network as it is observed from the throughput-delay curves, the influence of the GTW, is of no considerable importance.

ACKNOWLEDGEMENT

The authors would like to thank Dr. S. Koumbias for his continuous encouragement and helpful discussions. This work was carried out under the partial financial support of the EEC (ESPRIT project LION 169) as a subcontract with THOMSON-TITN, France.

REFERENCES

[1] Metcalf, R.M.: "Ethernet: Distributed packet switching for local computer networks", Communications of the ACM 19,7 (July 1976), pp.395-404.

[2] Kleinrock, L., Tobagi, F.A.: "Packet switching in Radio channels: Part I-Carrier Sense Multiple Access Modes and their throughput-delay characteristics", IEEE Trans. Commun., Vol. COM-23, pp.1400-1416, 1975.

[3] ESPRIT Pilot Project: "Local Wideband Communication System", Pilot Phase Report, Sept. 1984.

[4] Heidelberger, P., Lavenberg, S.S.:"Computer Performance Evaluation Methodology", IEEE Trans. Comp., Vol. C-33, pp.1195-1219,1984.

[5] Grogono, Peter: "Programming in Pascal", Addison-Wesley Series, 1979.

[6] "The ETHERNET, A Local Area Network", version 2.0, Digital Equipment Corporation, Intel Corporation, Xerox Corporation, November 1982.