

Sensitivity analysis of the load on an ethernet network

Citation for published version (APA):

Dekkers, A., & Oorschot, van, J. P. M. (1991). *Sensitivity analysis of the load on an ethernet network*. (Memorandum COSOR; Vol. 9141). Technische Universiteit Eindhoven.

Document status and date:

Published: 01/01/1991

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

TECHNISCHE UNIVERSITEIT EINDHOVEN
Faculteit Wiskunde en Informatica

Memorandum COSOR 91-41

Sensitivity analysis of the load on
an ethernet network

A. Dekkers
J. van Oorschot

Eindhoven University of Technology
Department of Mathematics and Computing Science
P.O. Box 513
5600 MB Eindhoven
The Netherlands

Eindhoven, December 1991
The Netherlands
ISSN 0926-4493

SENSITIVITY ANALYSIS OF THE LOAD ON AN ETHERNET NETWORK

A. Dekkers¹

Eindhoven University of Technology
Department of Mathematics and Computing Science
P.O. Box 513
5600 MB Eindhoven, The Netherlands
e-mail: wscoade@win.tue.nl

J. van Oorschot

Delft University of Technology
Department of Electrical Engineering
P.O. Box 5031
2600 GA Delft, The Netherlands
e-mail: JPMvOorschot@ET.TUdelft.NL

Abstract

In this paper a sensitivity analysis of the load on an ethernet network is presented. Aim is to determine which parameters that characterize the workload have an influence on the performance of an ethernet segment and to get an idea about the magnitude of this influence. The performance of the network is indicated by the waiting time and the access delay. The performance is measured using a simulation model. The simulation model is a very accurate description of a real ethernet implementation and has been validated with the results of real life measurements (Van Oorschot and Dekkers [9], Thio [6]). The main advantages of simulation are the detailed information that can be obtained and the absence of expensive laboratory measurements.

The importance of this sensitivity analysis is to validate future modelling assumptions and to get a better insight in the behaviour of the ethernet protocol. From the research it can be concluded that the performance is sensitive not just to the load of the ethernet, but also to the distribution of the interarrival times, the possibility of batch arrivals and the fraction of traffic arriving in large messages of a number of packets. Therefore existing models (e.g. Tobagi and Hunt [11], Lam [5], Bertsekas and Gallager [2]) only yield a rough performance estimation of a real ethernet segment, because they assume Poisson arrivals and do not consider batches.

¹ The investigations were supported (in part) by the Foundation for Computer Science in the Netherlands (SION) with financial aid from the Netherlands Organization for the Advancement of Scientific Research (NWO).

1. Introduction

The performance of an information processing system depends on many characteristics of the system. Only some of these can be controlled by performance analysts, especially those concerning the workload, and the sensitivity of the performance with respect to these parameters can be studied. In this paper a sensitivity analysis of an ethernet segment is presented. Ethernet is a well-known and widely spread broadcast medium, which was first described by Metcalfe and Boggs [7] in 1976. In our research, one segment is considered without taking into account the connections with the outer world through gateways or bridges. Of course, an ethernet-bridge can be connected to a LAN segment but then its transmissions are considered as being generated by that ethernet-node. In a previous paper [9] we showed that a standard IEEE 802.3 ethernet segment [1] can be simulated very accurately. Here a sensitivity analysis is presented, based on this simulation. Final goal of this research is to derive a model, which contains all aspects that influence the performance of the ethernet type of network. This is important for two reasons in future research. Firstly, in searching simple, analytical expressions for the performance we can concentrate on this model, knowing that other parameters than the ones included have a negligible effect on the performance. The current models make assumptions about certain parameters (often mentioned are Tobagi and Hunt [11], Lam [5], Bertsekas and Gallager [2]). Here we study whether these assumptions are too restrictive for the models or that these models are also good approximations if the assumptions are not fulfilled. Secondly, in measuring an extended situation we know with this information which input data is of importance and should be varied, and which input data influence the performance only slightly and can therefore be chosen freely. Other simulation studies on ethernet in which one aspect of the configuration is highlighted are Gonsalves and Tobagi [3] (the distance between stations) and Nutt and Bayer [8] (combined voice data load).

The main question is, which parameters have a significant influence on the performance of a segment and which parameters can be ignored. The performance is given by two parameters, the waiting time of a packet and the access delay of a packet. The waiting time is the time from the moment a packet arrives in the system, until its broadcast without collision is started. This time indicates how long it takes for a packet to be transmitted. The access delay is the interval from the moment a packet is in service position, i.e. first in queue, and ready to be transmitted until its broadcast without collision is started. This time indicates how many machines are contending for the right to transmit a packet. If a packet arrives on a busy machine its waiting time is larger than zero. But its access delay may be near zero, if the net is empty on the moment the packet becomes the first one in a queue. We anticipate that the two major parameters for the performance are the average packet length and the average interarrival time between two arrivals. Directly derived from these parameters is the utilization of the net, i.e. the fraction of the time the net is used for successful transmissions. To test the influence of another parameter we vary this parameter and compare the waiting time and access delay with a "standard" situation for several utilizations. If the average packet length is measured on a real ethernet implementation this yields two classes of packets. Most packets have either a short length close to the minimum packet length (between 64 and 200 bytes) or are of the maximum length (1514 bytes) and only a few packets have a length somewhere in between. In our measurements, only the fraction of small packets compared to large ones varies. In a real implementation the large packets originate from a message in which a lot of data has to be transferred. Therefore we consider two kinds of messages, "small" messages consisting of 1 small packet or "large" messages consisting of a number of large packets. Note the importance of the difference between messages and packets, which is made throughout the entire paper. The interarrival time distribution is in practice

different for each application. This holds for the expectation as well as for higher moments. In our simulation we found that the performance of an ethernet segment started to behave very poorly at a utilization of approximately 60%, a well known boundary from literature (e.g. Tanenbaum [10] p. 148, a mixture of long and small packets). Therefore we simulated up to a utilization of 70% to show this effect.

In this paper first a standard situation is described in Section 2. Section 3 depicts a rough model for an ethernet segment. In Section 4 the number of machines connected to the ethernet segment is discussed. Section 5 is dedicated to the coefficient of variation for the interarrival times. In Section 6 the effect of the fraction of small and large packets is given. The influence of the batchsize on the performance is discussed in Section 7. Section 8 contains some conclusions and suggestions for further research.

2. The "standard" situation

In this section the workload characterizations are described, which we consider typical for a real ethernet implementation. Of course, in practice the value of several parameters may vary. But with the given utilization the situation looks like the one described here. The ethernet we consider is a standard IEEE 802.3 ethernet segment [1] with a capacity of 10 Mbit/sec or equivalently a net speed of 0.0008 msec/byte. The collision recovery is handled with the truncated binary exponential backoff mechanism, the jam size is 32 bytes and the interframe gap is 0.096 msec. In Van Oorschot and Dekkers [9] we showed that the time T to transmit a packet can be represented by a linear function of the packetlength:

$$T(\text{packetlength}) = EC + \text{netspeed} \cdot \text{packetlength} \quad (1)$$

with the *packetlength* in bytes and *EC* a constant which depends on the speed and architecture of the ethernet card. Here we take *EC* as 0.03 msec, a representative number [9]. Essential is that *EC* is larger than the interframe gap, preventing one machine from sending several packets in a row, while other machines wait for the opportunity to transmit. The machines in the simulation all have an infinite buffer on their ethernet card. This corresponds in a real system to a buffer with several places and a main cpu, which can produce packets faster than the network card can transmit them. For a normal PC and packets of 1514 bytes this is true (Van Oorschot and Dekkers [9]). In the ethernet protocol packets which collided 16 times are removed from the ethernet card to prevent a system overload. This happens only seldom, but the access delay and waiting time of these packets until removal are included in our measurements.

We now describe a situation which we consider as representative for a real system. It is the "standard" situation in our further simulations and we vary one parameter at a time from this situation. In the **standard situation** there are eight machines connected to the ethernet segment, each machine generates the same workload. The traffic consists of small messages of one packet of 100 bytes and of large messages with 100 packets of 1514 bytes (here for simplicity the last packet also has the maximum length, in practice the last one may have a smaller length). As stated before, a message consists either of one packet or of 100 packets. A machine generates both kind of messages independently from one another, but transmits them via one ethernet card.

In measurements on the ethernet of the campus of the Delft University of Technology (Thio [6]) the number of small packets was three times as high as the number of large packets. Thus the ratio of large and small messages is 1 to 300. Here we use the same relation. This results in small average interarrival times for small messages, corresponding to e.g. a typing user and larger average interarrival times for large messages, corresponding to e.g. a file transfer. Note that, although there are less large packets than small packets (ratio 1 to 3), the large packets use more of the capacity of the net (ratio $1 \cdot 1514/3 \cdot 100 = 1514/300$). Important for the performance is the utilization of the net and we vary it from 0.1 up to 0.7 in our simulation. This implies the values of the average interarrival times of small and large messages to satisfy a specific utilization. At one machine there are two independent Poisson arrival streams: a stream of small messages, consisting of one packet, and a stream of large messages, consisting of 100 packets. If we vary one parameter, one of the values we take into account represents the value of this parameter in the standard situation. Therefore we did not plot the standard situation separately, it will be depicted in every figure by the dashed line with circles.

3. A rough model

Here we will give a model for the waiting time at an ethernet segment that will include the major delay effects. Especially the backoff mechanism makes an ethernet segment very hard to model. Therefore no model for the access delay is included. The contention solving mechanism results in a too complex system of idle machines, waiting machines and machines with an on-going backoff period to be modelled easily. A natural model for the waiting time at an ethernet segment seems to be a 1-limited polling model. The system can be represented by one server (the ethernet cable) and several queues (the machines) with jobs (packets) that have to be served. The server serves only one packet at a queue and then goes to another queue, this corresponds to the 1-limited discipline. There are, however, several difficulties in analysing a polling model of an ethernet segment; there are two distinct classes of jobs, there will be some loss of capacity due to collisions and backoff intervals and the service order is according to a random mechanism. Note that the net is occupied mostly by large packets. A model for the average waiting time for an average large packet is now to consider only large packets arriving in batches. Now, according to Groenendijk[4] the waiting time W_l for a large packet is

$$W_l = \frac{\beta_l}{2(1 - \rho)} (\rho + K - 1), \quad (2)$$

with β_l the transmission time of a large packet (netspeed times packetlength), ρ the utilization and K the batchsize. For the transmission time as well as the utilization the obligatory gaptime after a transmission is considered to be part of that transmission. This approximation formula can be adapted for small packets by Little's Law and some heuristic arguments. By Little's Law the number of large packets L_l machine equals

$$L_l = \lambda_l (W_l + \beta_l) \quad (3)$$

with λ_l the arrival intensity of large packets. An average large packet thus has to wait $\frac{1}{2}(K - 1)$ rounds (packets of its own batch in front of it) plus L_l rounds (the queue in front of it) minus $\frac{1}{2}\rho$ round (correction for on-going round) of the server before it is served. Hence, the cycle time C is

$$C = \frac{W_l}{L_l + \frac{1}{2}(K - 1) - \frac{1}{2}\rho} \quad (4)$$

The waiting time for a small packet W_s , arriving in equilibrium, is approximately

$$W_s = C \left(L_l - \frac{1}{2}\rho \right), \quad (5)$$

because it does not arrive in a batch and thus only has to wait on the queue present at its arrival moment.

The major drawback of the model above is that a fair random access mechanism is assumed, while in practice large messages tend to dominate the net. Consider the following scenario: after the transmission of a large packet from machine A, this machine waits while several other machines contend for access to the net. If after EC msec. the contention is not solved and all machines are waiting for their backoff time to expire, machine A starts transmitting again. At the end of its transmission, this same procedure can happen. The order of service is thus often not a random visit to a queue with packets, but frequently a sequence of visits to one machine with packets of a large message, especially at heavy load. If all packets of one message were served consequentially, an M/G/1 model would be a good first order approximation. Then

$$W_s = \frac{\rho}{1 - \rho} R \quad (6)$$

and

$$W_l = \frac{\rho}{1 - \rho} R + \frac{1}{2}(K - 1)(\beta_l + EC) \quad (7)$$

with R the residual service time of a message.

In practice a certain number of packets of one message are served consequentially, thus only a part of a message, i.e. nor a polling model nor a M/G/1 model is correct, but a mixture of both these models. It is obvious that the domination of the net by large messages depends on the utilization, for the probability on several contending machines is higher with a higher utilization. As a rough approximation for the waiting time in an ethernet segment we use a hybrid model, in which the waiting time is simply ρ times the waiting time in a M/G/1 model and $(1 - \rho)$ times the waiting time in a polling model. For a more accurate model the contention protocol, which is very complex, should be examined carefully. Note that also the loss of capacity due to collisions should then be taken into account. This remains an interesting research topic. Another topic of research should be if this protocol could be optimized by preventing one machine from dominating the net. In the next figures the results of the simulation and of the model for the standard situation are given.

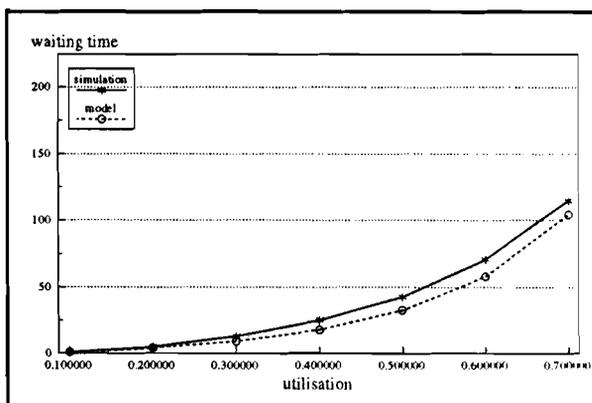


Fig. 1: Waiting time of small packets.

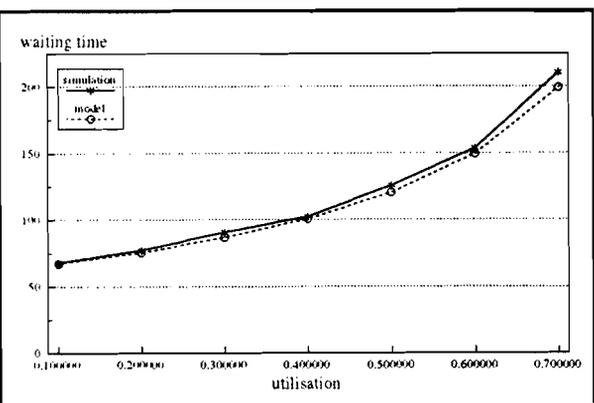


Fig. 2: Waiting time of large packets.

4. Number of machines

A question in the simulation of an ethernet segment is how many machines should be connected to the net and generate traffic. In a simulation it is easy to increase the number of machines almost unbounded, but for actual measurements there is a restriction. We are therefore interested in a representative number of machines. A representative number of machines would be the minimal number, for which the addition of another machine does not influence the waiting time, given an amount of traffic has to be broadcasted. In the next figures we see the influence of the number of machines on the waiting time and on the access delay for small messages as well as for large messages. For the waiting time of large messages we plotted the average waiting time packet for packets out of a large message and the total waiting time until the transmission of the last packet of a large message is started successfully. The traffic is generated by 2, 4, 8 or 16 machines. The rest of the situation is equal as described for the standard situation, except for the average interarrival times. These are adjusted to meet the required utilization.

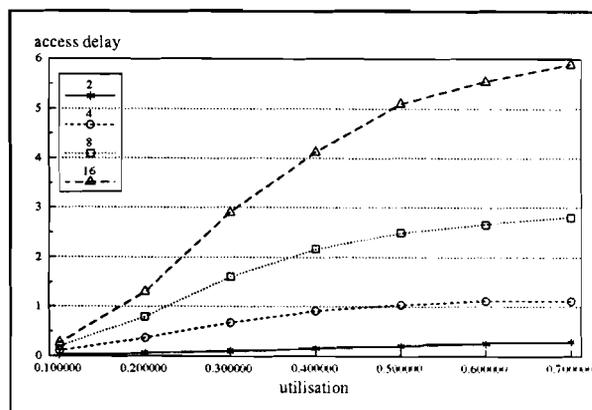


Fig. 3: Access delay of small packets for several numbers of machines.

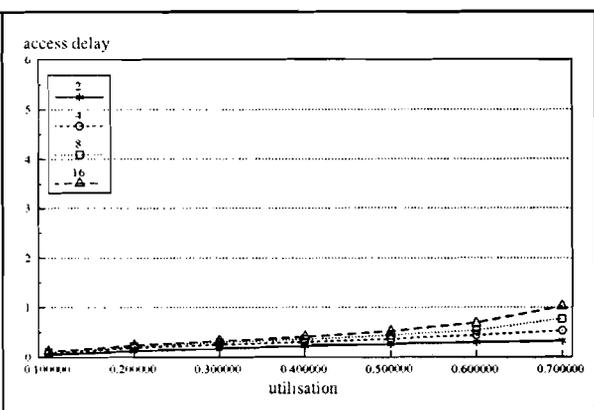


Fig. 4: Access delay of large packets for several numbers of machines.

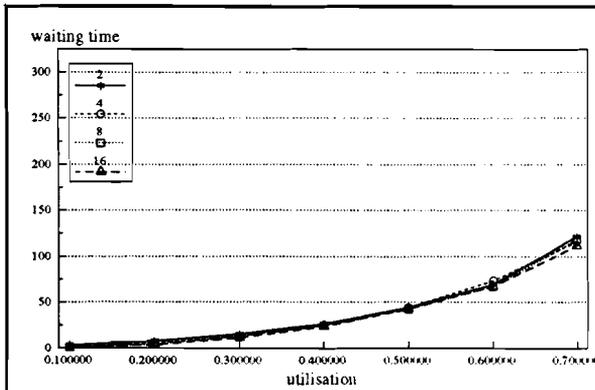


Fig. 5: Waiting time of small packets for several numbers of machines.

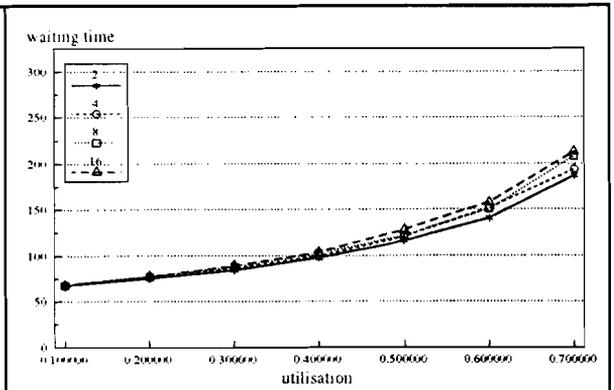


Fig. 6: Average waiting time of large packets for several numbers of machines.

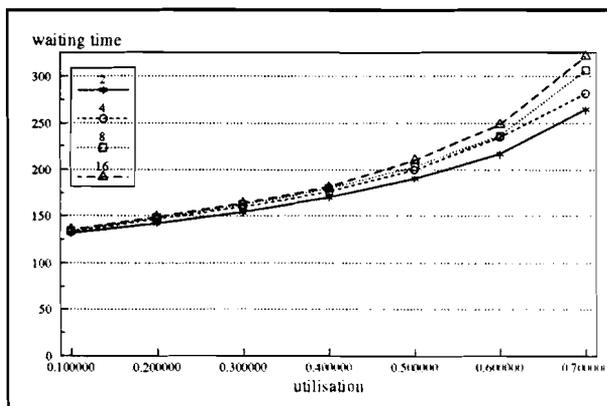


Fig. 7: Total waiting time of large messages for several numbers of machines.

From the figures above it is clear that at about 60% utilization the performance of the ethernet decreases dramatically. The waiting time now grows very rapidly, but the access delay still shows only a moderate growth. The growth of the access delay was expected, because with less machines there is less contention, and thus less collisions. The access delay of large packets differs from the access delay of small packets, because large packets can dominate the net as described in the previous section. Note that if one large message dominates the net, other large messages have to wait too. With an increasing utilization, the probability that there are two large messages at the same time increases. The access delay for large packets therefore grows rapidly with the utilization, for the access delay is mainly determined by large packets of another machine. For small packets it is important whether there are large messages in the system or not, more than the question how many large messages there are. The access delay for small packets therefore grows less rapidly. The number of machines has a small influence on the waiting time. This was expected by the outcomes of our model. For small messages there is an advantage of more machines in the polling model. This advantage is diminished by the unfair

access mechanism and by the fact that with more machines there are more collisions and thus more capacity loss. For large messages the number of machines does not occur in our model. But there are differences in the simulation because of the loss of capacity and because the effects for small messages should in an accurate model also influence the waiting time of large packets.

In the remainder of this paper we generate the traffic with 8 machines. This choice is rather arbitrarily. We think that 8 machines is a realistic number and the effect of adding a few machines is small as can be seen in the simulation and in our model.

5. Distribution of interarrival times

In the previous paragraph there were two kind of arrivals at one machine, the arrival of large messages and the arrival of small messages. Both types of messages arrived independently to one another according to a Poisson process. In this section we show the effect of varying the interarrival times for one type. The large messages and the small ones still arrive independently from one another. The interarrival times are in this section not exponentially distributed, but drawn from Γ -distributions with coefficients of variation of $\frac{1}{2}$ (Erlang 4), $\frac{1}{2}\sqrt{2}$ (Erlang 2), 1 (exponential), $\sqrt{2}$ and 2. For each of these values the waiting time and access delay are plotted as a function of the utilization for small and large messages. Again the total and average waiting time for large messages are given. The rest of the situation is equal to the standard situation.

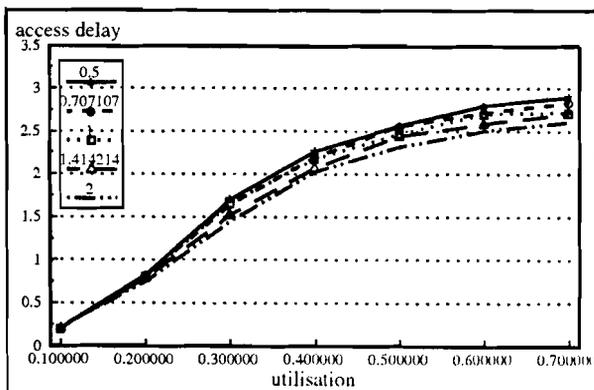


Fig. 8: Access delay of small packets for several coefficients of variation.

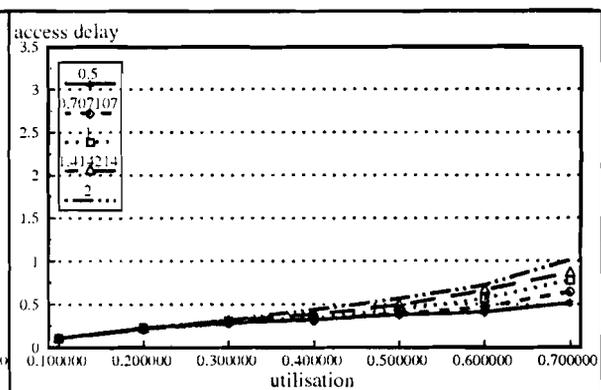


Fig. 9: Access delay of large packets for several coefficients of variation.

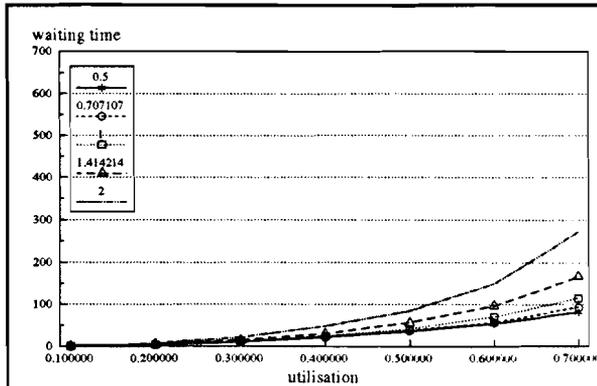


Fig. 10: Waiting time of small packets for several coefficients of variation.

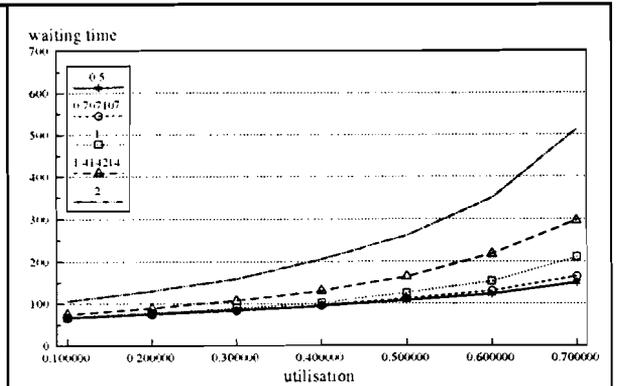


Fig. 11: Average waiting time of large packets for several coefficients of variation.

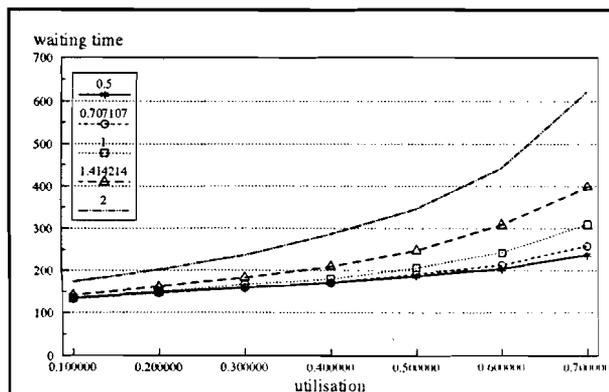


Fig. 12: Total waiting time of large messages for several coefficients of variation.

The waiting time of small as well as large packets is larger if the coefficient of variation is larger. This is the well known result that more randomness increases the waiting time. Therefore the arrival process should be included in an accurate model of an ethernet segment, because it is of importance for the performance. For the access delay the situation is very complex. Besides the effect above, there is an opposite effect. A high coefficient of variation generates bursty traffic on a machine. But being a packet in a burst does not influence the access delay. The access delay is a measurement of bursts on the other machines.

6. Distribution of the packetlength

In measurements at the Delft University of Technology the packetlength distribution was 75% small packets and 25% large packets. With a batchsize of 100 packets for large messages this is equivalent to a ratio of 300 small messages to 1 large message. The measurements are of a university environment and may be different in the banking or business world, or even at other

universities. In all situations however you see mainly small (between 64 and 200 bytes) and large packets (1514 bytes), only non-optimal protocols or incidental occurrences yield other packet lengths. Therefore we only consider packets of 100 and 1514 bytes. In this section the percentage of small packets of the total number of packets is varied from 55% to 95%. This means that the ratio of small messages and large messages varies from 122.22:1 to 1900:1. In the simulation the rest of the situation is equal to the standard situation. We plotted the waiting time and access delay as a function of the utilization for five different fractions of small packets: 0.55, 0.65, 0.75, 0.85 and 0.95:

Percentage small packets	ratio small/large packets	ratio small/large messages
55%	$55/45=1.22/1$	122.2/1
65%	$65/35=1.86/1$	185.7/1
75%	$75/25=3/1$	300/1
85%	$85/15=5.67/1$	566.7/1
95%	$95/5=19/1$	1900/1

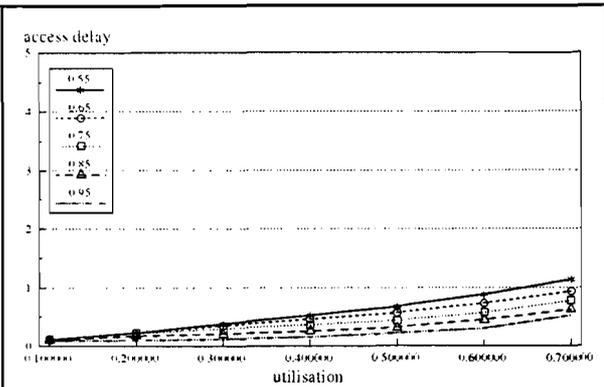
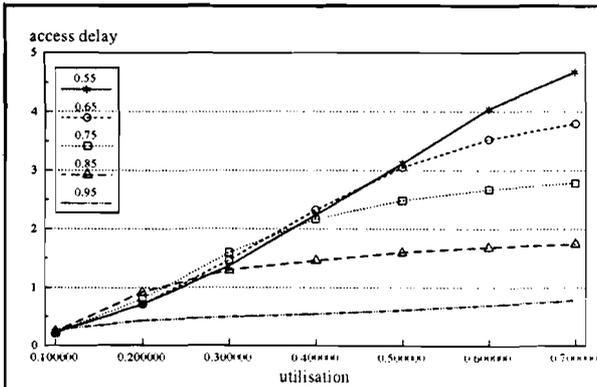


Fig. 13: Access delay of small packets for several fractions of small packets.

Fig. 14: Access delay of large packets for several fractions of small packets.

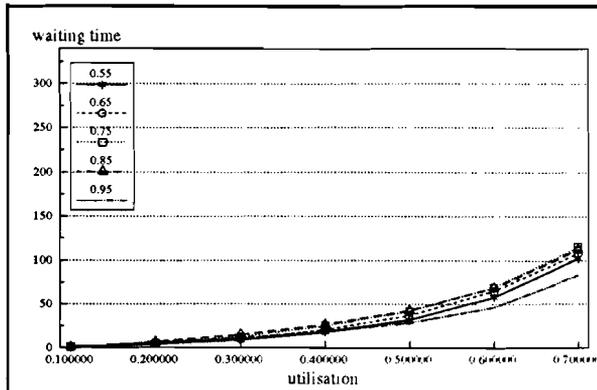


Fig. 15: Waiting time of small packets for several fractions of small packets.

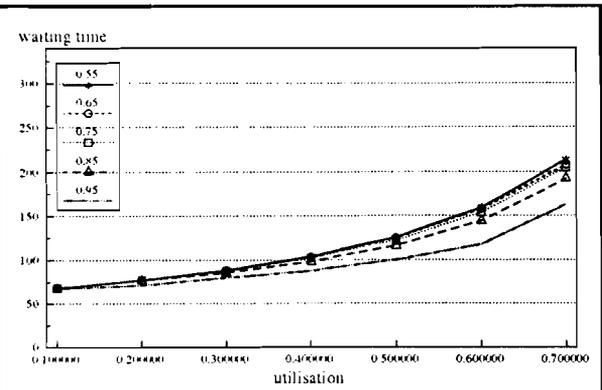


Fig. 16: Average waiting time of large packets for several fractions of small packets.

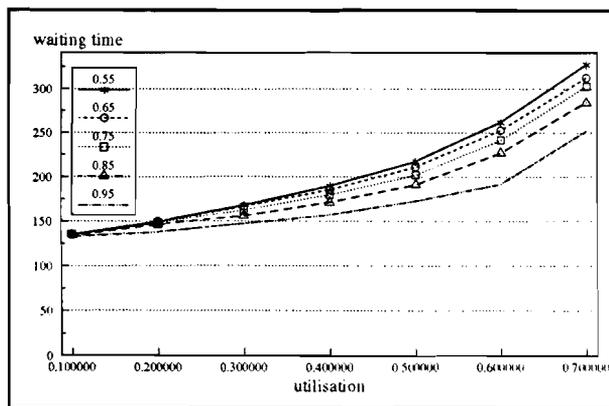


Fig. 17: Total waiting time of large messages for several fractions of small packets.

The difference in waiting time for large packets can be declared by the M/G/1 part of our model. If the fraction of large packets is increased, also the residual service time, and thus the waiting time, increases. Another important effect for small packets is that the absolute number of packets decreases if the fraction of large packets increases. This results in less collisions and thus in less loss of capacity. Both effects work in opposite directions, and which one is most important is not clear beforehand. In our rough model the number of small packets is not included. For the access delay again there are two opposite effects, which are both given above. More study should yield very accurate models for the access delay to be able to declare the figures above.

7. Batchsize

Up to now only fixed size batches of 100 packets are considered. We anticipate that it is of importance whether or not there are batches of considerable length. This means that there are (short) periods with a very high load and periods with a relatively low load. Of course the

batchsize influences the waiting time of large messages, but for the access delay of both types of messages we expect no influence if the batchsize is not too small. Here we take batches of 1, 10, 100 and 200 packets. The interarrival times of messages are adapted to result for a certain utilization in an equal throughput of large packets in each of these cases. The rest of the parameters was set as in the standard situation. This yields the next figures for the waiting time and access delay as a function of the utilization.

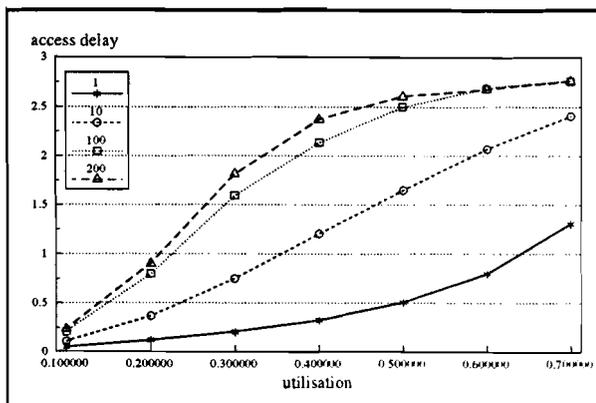


Fig. 18: Access delay of small packets for several batchsizes.

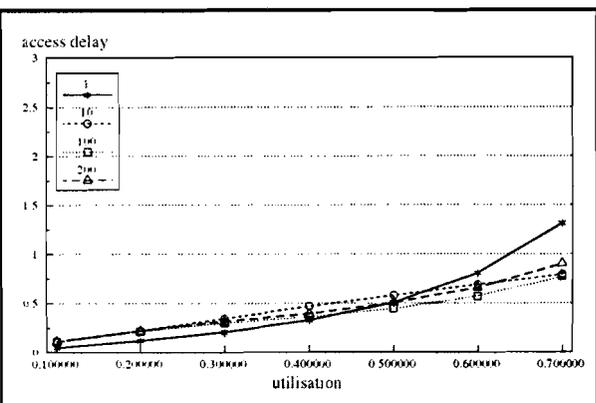


Fig. 19: Access delay of large packets for several batchsizes.

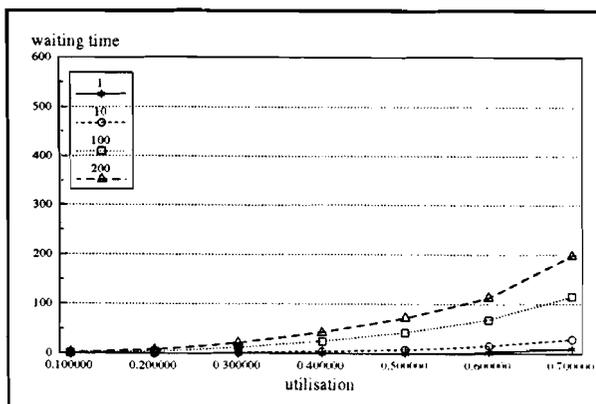


Fig. 20: Waiting time of small packets for several batchsizes.

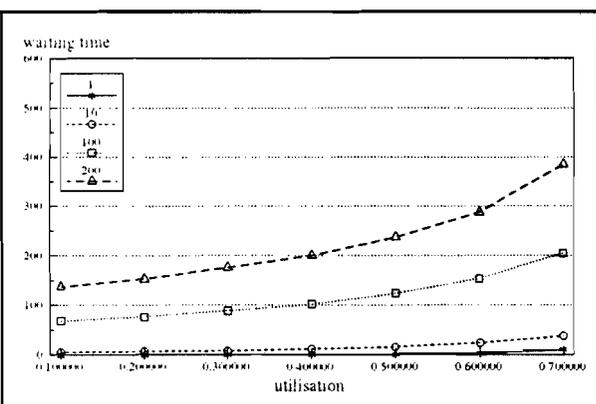


Fig. 21: Average waiting time of large packets for several batchsizes.

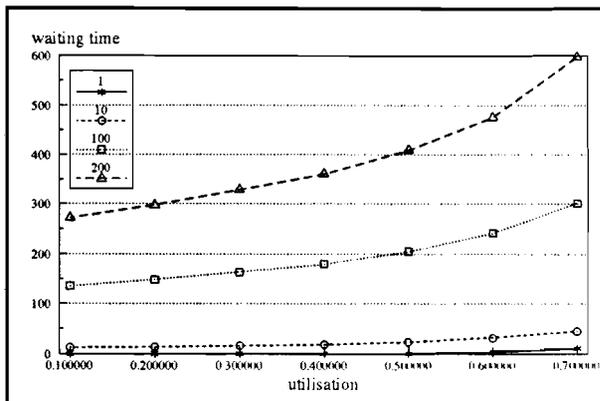


Fig. 22: Total waiting time of large messages for several batchsizes.

Our model is only valid if there really exists batches of a considerable size. Then it can be used to yield a reasonable approximation. For small batches (from 1 or 10) the situation is essentially different and another approach will be needed. The main difference is that no machine will be able to dominate the net during a longer interval of time. The waiting time for small packets depends mainly on the number of large packets in front of it. With large batches the probability on a very large queue in front of it increases. Therefore the waiting time of small packets is sensitive to the batchsize of large messages. The fundamental difference between a situation with or without batches is clearly shown by figure 19. The line for batches of size 1 is equal to the line for small packets and completely different from the other lines. This once again shows that the access delay is a difficult but important performance measure.

8. Conclusions and suggestions for further research

In this paper the influence of the workload characterization on the performance of an IEEE 802.3 ethernet segment is discussed. As expected, the performance of such a segment is mainly determined by the utilization of the net. The influence of the other parameters is tested using a standard situation and varying only the tested parameter.

We introduced a very rough model for the waiting time. This model is based on a combination of a polling model and a M/G/1 model. Major questions about this model are how to mix these two models appropriately and how to include collisions and attendant capacity loss. Furthermore there is by our knowledge no polling model that is able to deal with two completely different types of customer on one machine nor a polling model for a random moving server. To solve all these types of problems will require thorough research. The latter problems are theoretical, while the major ones can only be solved by approximations and a very good understanding of how the complex ethernet protocol is defined. The access mechanism is very complicated and therefore very hard to model. A complete understanding of this protocol and its practical behaviour is needed to be able to model it. The problems are caused by the backoff mechanism and by the loss of capacity due to collisions.

From the results we concluded that the influence of the number of machines is not very large, we think that 8 machines is a very reasonable choice.

The other parameters we tested have a larger influence on the performance and if one is interested in accurate results these factors should be included in the models. The existence of batches is essential for the functioning of the protocol. One large message can dominate the net. It is an interesting question if this domination can be avoided easily and what will be gained by this alteration of the ethernet protocol.

We can conclude that the performance of an ethernet segment depends heavily on kind of traffic it has to transmit. Above 60% utilization the performance starts to degrade rapidly. The ethernet protocol remains a difficult, but challenging, subject for performance analysts. Due to the hard protocol it requires much effort to make a proper model and due to the popularity of ethernet this is of importance. With this paper we have indicated why it is hard to model ethernet and we posed questions which may help in further research.

9. References

- [1] *IEEE Standards for Local Area Networks: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications* IEEE, New York, 1985.
- [2] BERTSEKAS, D. AND GALLAGER, R., *Data Networks*, Prentice-Hall, Englewood Cliffs, New Jersey, 1987.
- [3] GONSALVES, T.A. AND TOBAGI, F.A., On the Performance Effects of Station Locations and Access Protocol Parameters in Ethernet Networks, *IEEE Trans. on Comm.*, V 36, pp. 441-449, 1988.
- [4] GROENENDIJK, W.P., *Conservation Laws in Polling Systems*, Ph.D. Thesis, Dep. of Math. and Comput.Sci., University of Utrecht, Utrecht, 1990.
- [5] LAM, S.S., A Carrier Sense Multiple Access Protocol for Local Networks, *Computer Networks*, V 4, pp. 21-32, 1980.
- [6] THIO, L., *Measuring the workload on the Delft University of Technology*, Delft University of Technology, Internal report, Delft, 1990.
- [7] METCALFE, R.M. AND BOGGS, D.R., Ethernet: Distributed Packet Switching for Local Computer Networks, *Comm. of the ACM*, V 19, July 1976, pp. 395-404.
- [8] NUTT, G.J. AND BAYER, D.L., Performance of CSMA/CD Networks Under Combined Voice and Data Loads, *IEEE Trans. on Comm.*, V 30, pp. 6-11, 1982.
- [9] OORSCHOT, J. VAN AND DEKKERS, A., *Measuring and Simulating an 802.3 CSMA/CD LAN*, Eindhoven University of Technology, Department of Mathematics and Computing Science, Memorandum COSOR 90-26, Eindhoven, 1989.
- [10] TANENBAUM, A.S., *Computer Networks*, Prentice-Hall, Englewood Cliffs, New Jersey 1988, second edition.
- [11] TOBAGI, F.A. AND HUNT, V.B., Performance Analysis of Carrier Sense Multiple Access with Collision Detection, *Computer Networks*, V 4, pp. 245-259, 1980.

EINDHOVEN UNIVERSITY OF TECHNOLOGY
Department of Mathematics and Computing Science
PROBABILITY THEORY, STATISTICS, OPERATIONS RESEARCH
AND SYSTEMS THEORY
P.O. Box 513
5600 MB Eindhoven, The Netherlands

Secretariate: Dommelbuilding 0.03
Telephone : 040-473130

-List of COSOR-memoranda - 1991

<u>Number</u>	<u>Month</u>	<u>Author</u>	<u>Title</u>
91-01	January	M.W.I. van Kraaij W.Z. Venema J. Wessels	The construction of a strategy for manpower planning problems.
91-02	January	M.W.I. van Kraaij W.Z. Venema J. Wessels	Support for problem formulation and evaluation in manpower planning problems.
91-03	January	M.W.P. Savelsbergh	The vehicle routing problem with time windows: minimizing route duration.
91-04	January	M.W.I. van Kraaij	Some considerations concerning the problem interpreter of the new manpower planning system formasy.
91-05	February	G.L. Nemhauser M.W.P. Savelsbergh	A cutting plane algorithm for the single machine scheduling problem with release times.
91-06	March	R.J.G. Wilms	Properties of Fourier-Stieltjes sequences of distribution with support in $[0,1)$.
91-07	March	F. Coolen R. Dekker A. Smit	Analysis of a two-phase inspection model with competing risks.
91-08	April	P.J. Zwietering E.H.L. Aarts J. Wessels	The Design and Complexity of Exact Multi-Layered Perceptrons.
91-09	May	P.J. Zwietering E.H.L. Aarts J. Wessels	The Classification Capabilities of Exact Two-Layered Peceptrons.
91-10	May	P.J. Zwietering E.H.L. Aarts J. Wessels	Sorting With A Neural Net.
91-11	May	F. Coolen	On some misconceptions about subjective probability and Bayesian inference.

COSOR-MEMORANDA (2)

91-12	May	P. van der Laan	Two-stage selection procedures with attention to screening.
91-13	May	I.J.B.F. Adan G.J. van Houtum J. Wessels W.H.M. Zijm	A compensation procedure for multiprogramming queues.
91-14	June	J. Korst E. Aarts J.K. Lenstra J. Wessels	Periodic assignment and graph colouring.
91-15	July	P.J. Zwietering M.J.A.L. van Kraaij E.H.L. Aarts J. Wessels	Neural Networks and Production Planning.
91-16	July	P. Deheuvels J.H.J. Einmahl	Approximations and Two-Sample Tests Based on P - P and Q - Q Plots of the Kaplan-Meier Estimators of Lifetime Distributions.
91-17	August	M.W.P. Savelsbergh G.C. Sigismondi G.L. Nemhauser	Functional description of MINTO, a Mixed INTEger Optimizer.
91-18	August	M.W.P. Savelsbergh G.C. Sigismondi G.L. Nemhauser	MINTO, a Mixed INTEger Optimizer.
91-19	August	P. van der Laan	The efficiency of subset selection of an almost best treatment.
91-20	September	P. van der Laan	Subset selection for an -best population: efficiency results.
91-21	September	E. Levner A.S. Nemirovsky	A network flow algorithm for just-in-time project scheduling.
91-22	September	R.J.M. Vaessens E.H.L. Aarts J.H. van Lint	Genetic Algorithms in Coding Theory - A Table for $A_3(n,d)$.
91-23	September	P. van der Laan	Distribution theory for selection from logistic populations.

COSOR-MEMORANDA (3)

91-24	October	I.J.B.F. Adan J. Wessels W.H.M. Zijm	Matrix-geometric analysis of the shortest queue problem with threshold jockeying.
91-25	October	I.J.B.F. Adan J. Wessels W.H.M. Zijm	Analysing Multiprogramming Queues by Generating Functions.
91-26	October	E.E.M. van Berkum P.M. Upperman	D-optimal designs for an incomplete quadratic model.
91-27	October	R.P. Gilles P.H.M. Ruys S. Jilin	Quasi-Networks in Social Relational Systems.
91-28	October	I.J.B.F. Adan J. Wessels W.H.M. Zijm	A Compensation Approach for Two-dimensional Markov Processes.
91-29	November	J. Wessels	Tools for the Interfacing Between Dynamical Problems and Models withing Decision Support Systems.
91-30	November	G.L. Nemhauser M.W.P. Savelsbergh G.C. Sigismondi	Constraint Classification for Mixed Integer Programming Formulations.
91-31	November	J.Th.M. Wijnen	Taguchi Methods.
91-32	November	F.P.A. Coolen	The Theory of Imprecise Probabilities: Some Results for Distribution Functions, Densities, Hazard Rates and Hazard Functions.
91-33	December	S. van Hoesel A. Wagelmans	On the P-coverage problem on the real line.
91-34	December	S. van Hoesel A. Wagelmans	On setup cost reduction in the economic lot-sizing model without speculative motives.
91-35	December	S. van Hoesel A. Wagelmans	On the complexity of post-optimality analysis of 0/1 programs.
91-36	December	F.P.A. Coolen	Imprecise Conjugate Prior Densities for the One-Parameter Exponential Family of Distributions.

COSOR-MEMORANDA (4)

91-37	December	J. Wessels	Decision systems; the relation between problem specification and mathematical analysis.
91-38	December	J.J.A.M. Brands R.J.G. Wilms	On the asymptotically uniform distribution modulo 1 of extreme order statistics.
91-39	December	C.J. Speck J. van der Wal	The capacitated multi-echelon inventory system with serial structure: 1. The 'push ahead'-effect.
91-40	December	C.J. Speck J. van der Wal	The capacitated multi-echelon inventory system with serial structure: 2. An average cost approximation method.
91-41	December	A. Dekkers J. v. Oorschot	Sensitivity analysis of the load on an ethernet network.