National Library of Canada
Acquisitions and Bibliographic Services Branch
395 Wellington Street
Ottawa, Ontario
K1A 0N4

Bibliothèque nationale du Canada
Direction des acquisitions et des services bibliographiques
395, rue Wellington
Ottawa (Ontario)
K1A 0N4

Your file Votre référence

Our file Notre référence

NOTICE

The quality of this microform is heavily dependent upon the quality of the original thesis submitted for microfilming. Every effort has been made to ensure the highest quality of reproduction possible.

If pages are missing, contact the university which granted the degree.

Some pages may have indistinct print especially if the original pages were typed with a poor typewriter ribbon or if the university sent us an inferior photocopy.

Reproduction in full or in part of this microform is governed by the Canadian Copyright Act, R.S.C. 1970, c. C-30, and subsequent amendments.

AVIS

La qualité de cette microforme dépend grandement de la qualité de la thèse soumise au microfilmage. Nous avons tout fait pour assurer une qualité supérieure de reproduction.

S'il manque des pages, veuillez communiquer avec l'université qui a conféré le grade.

La qualité d'impression de certaines pages peut laisser à désirer, surtout si les pages originales ont été dactylographiées à l'aide d'un ruban usé ou si l'université nous a fait parvenir une photocopie de qualité inférieure.

La reproduction, même partielle, de cette microforme est soumise à la Loi canadienne sur le droit d'auteur, SRC 1970, c. C-30, et ses amendements subséquents.

Canada

Design and Validation of Multicasting in an XTP Simulator

Rosarinho D'Cruz

A Major Report

in

The Department

of

Computer Science

Presented in Partial Fulfilment of the Requirements
for the Degree of Master of Computer Science at
Concordia University
Montreal, Quebec, Canada

May 1992

© Rosarinho D'Cruz, 1992

National Library of Canada
Bibliothèque nationale du Canada

Acquisitions and Bibliographic Services Branch
Direction des acquisitions et des services bibliographiques

395 Wellington Street
Ottawa, Ontario
K1A 0N4

395, rue Wellington
Ottawa (Ontario)
K1A 0N4

Your file Votre référence

Our file Notre référence

The author has granted an irrevocable non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of his/her thesis by any means and in any form or format, making this thesis available to interested persons.

L'auteur a accordé une licence irrévocable et non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de sa thèse de quelque manière et sous quelque forme que ce soit pour mettre des exemplaires de cette thèse à la disposition des personnes intéressées.

The author retains ownership of the copyright in his/her thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without his/her permission.

L'auteur conserve la propriété du droit d'auteur qui protège sa thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

ISBN 0-315-00921-3

Canada

ABSTRACT

Design and Validation of Multicasting in an XTP Simulator

R. D'Cruz

This report presents a design and validation of Multicasting extensions to an existing Unicasting XTP simulator for a transport and network level protocol, the Xpress Transport Protocol (XTP).

The tool on which the XTP simulator is based, is the Local Area Network Simulation Facility (LANSF). Ethernet is simulated as the LAN environment.

Analysis of XTP Multicasting performance indicates that its performance is about half of the unicasting performance for small messages, and approaches unicasting performance for large messages. This is because multicasting has larger overhead than unicasting. In multicasting, for every message sent whose size is less than or equal to 1460 bytes, one extra control packet has to be transmitted. Two extra control packets are required for message sizes greater than 1460 bytes. This extra overhead reduces the multicasting performance capability.

ACKNOWLEDGEMENTS

My sincerest thanks to my supervisor, Dr. J.W. Atwood, whose advice, criticism and patience have contributed very much to the completion of this major report.

Being married, having three children, having a full-time job, and being a part-time graduate student was not easy. I would like to thank my wife and my children for their patience and bearing with me.

TABLE OF CONTENTS

TABLE OF CONTENTS (Cont'd)

LIST OF TABLES

LIST OF TABLES (Cont'd)

LIST OF TABLES (Cont'd)

LIST OF FIGURES

LIST OF FIGURES (Cont'd)

CHAPTER 1

INTRODUCTION

The rapid technology advancement of distributed systems and high speed communication media requires an efficient transport layer protocol for high media bandwidth and throughput. Among the many protocols being studied, the Xpress Transport Protocol (XTP) is one example.

LANSF (Local Area Network Simulation Facility), which is a Configurable System for Modelling Communication Networks, was developed by Gburzynki and Rudnicki [2] at the University of Alberta. LANSF is a simulation model for investigating the performance of communication networks.

An XTP Simulator, which was writte. n LANSF V1.4, for performance analysis of the Eth .net environment at Revision 3.3 of the XTP specification, was completed for the master's thesis by Jason Chen [1].

This major report presents a design and validation of the addition of multicasting capability to the XTP simulator at Revision 3.5 of the XTP specification. For reasons discussed in section 2.5.1, this report utilizes the LANSF V2.11e. This report consists of five sections.

Section 2 presents background on the XTP Unicasting simulation. This section basically summarizes the vi/tal points required for multicasting from section 5 of Chen's thesis [1].

Section 3 introduces XTP Multicasting, and describes the bucket algorithm, which is applied to multicast operation.

Section 4 presents the design of XTP Multicasting mode in the simulation program. This section also describes the simulation plan and discusses the simulation results.

Section 5, the last section, presents the conclusion with discussion of the results obtained from the XTP performance simulation. This section also proposes future enhancements to the XTP simulator.

CHAPTER 2

BACKGROUND ON UNICASTING SIMULATION

The thesis "Design and Validation of an XTP Simulator" [1] in unicasting mode was based on the XTP specification, Version 3.3. This XTP Simulator is the base that has been modified for the inclusion of multicasting. The XTP Simulator uses an Ethernet environment, within which a number of stations set up virtual circuit connections to each other, and transfer data using the XTP.

This report identifies the essential points in the thesis and the changes required for multicasting. The essential points derived from the simulation plan (chapter 5) of the thesis [1], are described in the following subsections.

2.1 Simulation Assumptions

The Simulation assumptions are that there are four simulated "chips", each running its own task.

a) The main processor chip runs all the XTP processes (XTP Writer process, XTP Reader process, XTP Sender process and XTP Receiver process).

b) The timer chip handles all the XTP timer processes.

c) The rate controller chip does all the separation insertion between out-going packets.

d) The Ethernet controller chip performs all the CSMA/CD operations:

- The data link layer service that XTP uses is the IEEE 802.2 class I service, which is a connectionless service.
- All channels (queues) between simulation processes have unbounded length.
- Message fragmentation is not limited.

2.2 The Link and Physical Level Simulated Environment

The Ethernet transmitter initially waits for the sending signal from the serializer process or the timer process. The Ethernet receiver signals the XTP receiver after it has received a packet from the link and has mapped the packet to the input queue.

2.3 The Design of XTP Simulator

There are nine XTP processes: initialization process responsible for setting up the virtual circuit simulation data structures; xtp_writer and xtp_reader processes serve as transport service accessing points (TSAP); xtp_sender and xtp_receiver processes perform the core operations of XTP; the serializer process simulates the main processor, which runs the xtp_writer, xtp_reader, xtp_sender and xtp_receiver processes; the timer process does timer interrupt operations for ctimer, wtimer and rtimer; the credit control and rate control processes are assumed to be on a different chip.

The timer process has four timers in XTP which are maintained by the sending site: Context life timer (CTIMER) which is set to sixty seconds, Wait reply timer (WTIMER) which set to twice the estimated round trip delay, rate control timer (RTIMER), and the credit control timer (CRTIMER) which is set to one sixtieth (1/60) of a second. In XTP V3.3, the rate control and credit control were based on 1/60 sec, which is not the case for V3.5 . Since the rate control features of XTP have not been used in this study, it makes no difference whether 1/60 sec is used or not. In this simulation CTIMER and WTIMER are turned off during the no-error simulation, since the overall performance will not be affected.

2.4 The Simulation

The XTP simulation was run on the SUN4/SUNOS4.0 platform.

2.4.1 LANSF Tunable Parameters

The four LANSF tunable parameters, which can vary the rate of client level requests and which are of interest are:

(1) The message length (bits).

(2) The message interarrival time (message arrival rate in time units).

(3) The number of stations in a network.

(4) The number of messages that the LANSF client should generate for a simulation run.

2.4.2 XTP Tunable Parameters

The seven XTP tunable parameters in the XTP simulation are:

(1) alloc (receiving buffer size in bytes).

(2) credit (maximum number of bytes a virtual circuit can send in a burst).

(3) separation (minimum packet spacing on a route basis).

(4) wtimer (wait time value).

(5) ctimer (context life timer value).

(6) copy_delay (delay caused by copying data from and to user space).

(7) checksum_delay (on the incoming and outgoing packets).

2.4.3 Measurements

The measurements obtained are actual throughput and the individual average message delay time, for different message sizes under various traffic loads. The actual throughput is calculated as follows:

The actual throughput of the simulation = Delivered Sequence [byte]/Last send completed [s]

The individual average message delay time is calculated as follows:

The average message delay of the simulation = ((Offered Load [B/s] / Actual Throughput [B/s]) * (Minimum Message Delay))

The first round of the simulations were run without any copy-checksum delay to ensure that the maximum effective throughput corresponded to the expected maximum throughput.

2.4.4 **Simulation Plan**

The Simulation plan was as follows:

- The number of stations used in the XTP frame relay simulation was two: one sender and one receiver. Only the case of a single virtual circuit was simulated, because it would not make any difference at the MAC layer.
- The XTP wtimer and ctimer were turned off, since they do not affect the performance in the no-error environment. The XTP rtimer was also turned off, since there is no gateway involved in the simulation.
- The message sizes used in the simulation are 6, 128, 1024 and 8192 bytes for the following reasons:
  - The 6 byte message is the minimum number of bytes which can represent the terminal accessing activity, and which corresponds with the XTP minimum packet length criterion.
  - The 128 byte message represents remote procedure call activity.
  - The 1024 byte message represents page fetch operations.
  - The 8192 byte message represents file transfer operations.

There are ten different offered loads (testdata0 ... testdata9 input files) used in the simulation of each message size. The offered load was calculated by partitioning the expected maximum throughput into ten equal intervals. This would produce a linear curve. Three more offered loads (testdata10 ... testdata12) were added to the simulation to determine the saturation of the XTP performance. They were calculated by adding 10%, 20%, and 30% of extra load to the maximum expected throughput. The maximum expected throughput is calculated as follows:

The maximum expected throughput

= (User data length / Total packets length) * Total bandwidth

The Simulation of each message size at each offered load interval was run with and without copy-checksum delay. The LANSF client generated 10,000 messages to XTP in each simulation run. To ensure consistency of the simulation results, each simulation was run three time with different random number generation seeds, and the average of the three results was used.

The unicast simulation results for the different message sizes are indicated in figures 1, 2, 3 and 4 of this report. For the 8192 byte message size, the messages were fragmented (by LANSF) into five 1468 byte packets and one final packet of 852 bytes. The saturation of the throughput curves approximates with the calculated expected throughput.

2.5 Problems Encountered

2.5.1 XTP Simulator Upgrade for SUNOS4.1

The major problem was to upgrade the Unicast XTP Simulator using LANSF V1.4 to run on the SUN4/SUNOS4.1 system platform.

While the upgrade of the XTP Simulator for multicasting was in progress, the SUN4 SUNOS4.0 was upgraded to SUNOS4.1. The Unicasting XTP Simulator failed to run. Since LANSF V1.4 was not supported any more, LANSF had to be upgraded to V2.11e. The protocol.c file and some of the LANSF files had to be modified, in order to get the simulator running.

2.5.2 Utilization of XTP packet size

1) XTP packet size for 6 byte messages

In Chen [1], in the Unicast XTP simulation for 6 byte messages, the minimum ethernet packet size was calculated as follows:

frame_info_length + min_packet_length

= 464 (58 bytes) + 368 (46 bytes)

= 832 bits (104 bytes).

The minimum ethernet packet size should have been as follows:

frame_info_length + min_packet_length

= 208 (26 bytes) + 368 (46 bytes)

= 576 bits (72 bytes).

This correction has been accounted for in the XTP V3.5 simulation.

2) Fragmentation of 8192 byte messages

In Chen [1], the 8192 byte message was fragmented into five 1442 byte parts and one 982 byte part, as data sizes in XTP packets. The maximum data size in an ethernet packet is 1500 bytes. Therefore, the largest XTP packet size is 1500 bytes. Hence, the data size is

1500 - 32 (header + trailer in XTP V3.3) = 1468 bytes. Therefore, there should have been five 1468 byte parts and one 852 byte part, as data sizes in the XTP packets. These corrected values have been indicated in section 2.4.4.

For the XTP V3.5 simulation, the header plus trailer is equal to 40 bytes. Hence, there are five 1460 byte parts and one 892 byte part, as data sizes in the XTP packets.

CHAPTER 3

XTP MULTICASTING

3.1 General

XTP Multicast is intended for media that define a broadcast or multicast facility, i.e., where one sender can broadcast the same data stream or datagram sequence to multiple receivers simultaneously (one-to-many), and only one of the many receivers need respond when requested. The multicast service delivers data to receivers that are active for the duration of the multicast transmission.

3.2 Multicast Syntax

Multicast packets obey the same syntax rules as non-multicast: Header, Trailer, Control, and Information Segments are identical. Multicast packets differ from the unicast ones in the following ways:

1) All multicast packets have the MULTI bit turned on, and non-multicast packets clear the bit.
2) Multicast packets utilize group addresses rather than point-to-point addresses.
3) Multicast packets always use SREQ, not DREQ.
4) SREQ appears only in control (CNTL) packets,

not in DATA packets.

All other options and flags defined for unicast transmission are defined for multicast.

3.3 Multicast Control Packet Processing

The multicast flow and error control is same as in unicast. The definition of how a receiver should respond to the various options is the same as in unicast, except for the following:

- In the multicast mode, a sender must positively associate returned control packets with past events if continuous output streaming is desired. This can only be accomplished by matching returned echo values with local sync values. This in turn requires that a multicast sender set SREQ only in outgoing CNTL packets, not in FIRST or DATA packets.
- CNTL packets sent by multicast receivers in response to received SREQs are multicast to the group, not unicast to the sender. This is necessary for the slotting and damping algorithms.
- Key exchange is not possible in multicast mode. Consequently, CNTL packets to the multicast group that contain return keys are those generated by the

multicast receivers. CNTL packets containing non-primed keys are generated by the multicast sender.

Like the unicast sender, the multicast sender obeys the rules for flow control and rate control. When multicasting to a large set of receivers, Damping and Slotting techniques are supported that can be used to limit the duplicate or redundant control messages sent by the receivers. Even error control is supported, using the go-back-n retransmission scheme. Selective retransmission is not supported. However, for this study, error control is not handled, because the objective is to investigate throughput limitations.

3.4 Bucket Algorithm

Control messages received in unicast mode are acted upon without delay: dseq is used to release output buffers, rseq is used to start go-back-n retransmission, and alloc is used to advance the output window. In multicast mode some additional processing and delay must be inserted between the arrival of control messages and when action is taken. The purpose of the delay is to allow a reasonable time period for a multiplicity of receivers to respond to the SREQ or to report errors. The bucket algorithm described below manages this time delay.

A bucket is defined whenever the sender transmits a control packet with SREQ. The bucket refers to all data packets, possibly none, transmitted since the end of the previous bucket. This includes both transmissions and retransmissions. Each bucket is identified by the sender's **sync** value, which is computed by incrementing the **sync** value by 1 for every transmitted data packet.

When control messages are received from receivers they can be related to a particular bucket by their **echo** values, since **echo** is the returned value for **sync.** Responses are accumulated for each bucket as follows:

1) minimum of received **alloc**
2) minimum of received **dseq**
3) minimum of received **rseq**
4) maximum of received computed **rtt.**

If there are B buckets, then after a certain time period the accumulated alloc, dseq, rseq, rtt values for the oldest bucket are acted upon and the oldest bucket becomes the newest bucket. The algorithm detailing how this can be done is parameterized with respect to the following four variables:

1) WT (same as WTIMER), the current wait timer interval.
2) N, an implementation-selected value, which

must be 1 or greater, determines the number of SREQ control messages that will be transmitted within WT time interval.

3) ST, switching time, determined as WT/(N+0.5).

4) E, the desired multicast error tolerance, determines the error tolerance, which is defined as the number consecutive lost or damaged error control packets from a receiver to a multicast sender that will be tolerated.

A small value for E reduces message traffic, but increases the total amount of buffer space needed to run a multicast stream continuously (i.e., without falling into stop-and-wait behaviour). If no more than E packets are lost, then the retransmission algorithm will recover by retransmitting the lost data. If E is exceeded, then the sender's ability to retransmit will be exceeded and the receiver will never be able to recapture the missing data. Receivers must have methods for "dropping" out of multicast conversations when this occurs.

The details of the bucket algorithm can be stated as follows:

1) There are a total of N+E+1 buckets for an active multicast output context.

2) A control SREQ message is transmitted every ST seconds by the multicast sender. This is done even if no data is transmitted.

3) A new bucket is created for each SREQ.

4) Receivers multicast control responses to SREQ.

5) Incoming control messages at the multicast sender are associated with a bucket by matching returned **echo** value with the **sync** value on each bucket.

6) The bucket variables **alloc, dseq, rseq, rtt** are updated from the control packet, if the new values are less than **alloc, dseq, rseq**, or greater than **rtt**.

7) After N+E+1 buckets have been created, each new ST interval will relabel the "oldest" bucket as the "newest" bucket and the remaining buckets "age" one ST time unit. When this happens, the accumulated bucket variables are acted upon by the multicast transmitter. The values that are used are minimum **alloc, dseq, rseq,** over all buckets, and maximum **rtt** over all buckets.

The total amount of buffer space needed to operate the bucket algorithm continuously is:

Space = (N+E+1) * (WT/(N+0.5)) * F

= K * WT * F

where K = (N+E+1)/(N+0.5)

WT is the current wait timer (derived from rtt).

For increasing values of N, K approaches 1 which gives the minimum buffer space with the maximum number of buckets. If the number of buckets is reduced by setting N=1, then K=(E+2)/1.5 which leads to maximum buffer requirement. At both extremes for K the error tolerance E, is the same.

3.5 Multicasting Design

The XTP simulator has been modified to incorporate multicasting as defined in the XTP specification, Version 3.5. The bucket algorithm is applied in this design. The major difference between the unicasting and multicasting simulations is in the transmission of packets, where the overhead plays a role in the reduction of the performance. In multicast mode, SREQ must only appear in the control packets, and not in the DATA packets as it is in the unicast mode.

The design changes, which are applied, basically affects the xtp sender process and the xtp receiver process. The top level blockdiagram, which is given below, shows these processes. The modified functions are indicated modified in brackets. The bucket_algo function is added. The test files are also modified.

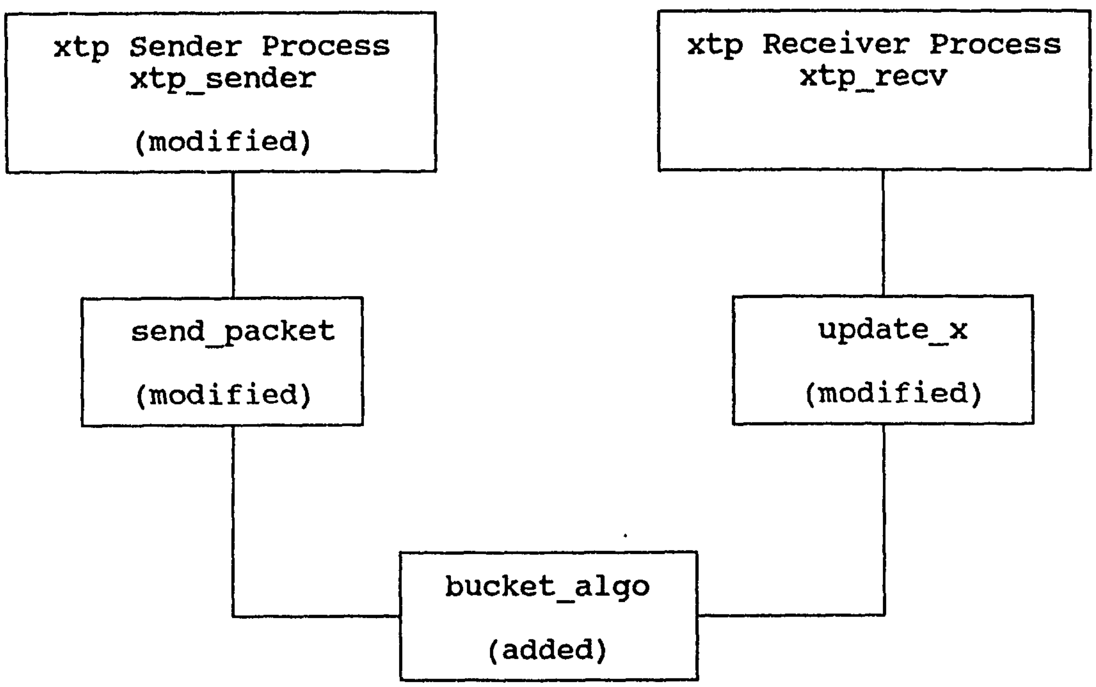

The sequence of packet transmission is as follows:

Unicasting:

```
Transmitter     ------->     Data Packet

                ------->

                   .

                   .

                   .

                ------->

Transmitter     ------->     Data Packet

                             with SREQ

CNTL Packet     <-------     Receiver

Transmitter     ——————>      Data Packet

                             with SREQ and EOM

CNTL Packet     <——————      Receiver
```

Note: The dashed lines are for message sizes greater than 1460 bytes.

Multicasting:

Transmitter	————————>	Data Packet
	-------->	
	.	
	.	
	.	
	-------->	
Transmitter	-------->	CNTL Packet
		with SREQ
CNTL Packet	<--------	Receiver
Transmitter	-------->	Data Packet
	-------->	
	.	
	.	
	.	
	-------->	
Transmitter	————————>	CNTL Packet
		with EOM and SREQ
CNTL Packet	<————————	Receiver

Note: The dashed lines are for message sizes greater than 1460 bytes.

CHAPTER 4

MULTICASTING SIMULATION

The same analogy of the XTP Unicasting simulation which was carried out in section 5.4 by Chen [1], has been summarized in section 2 of this report and applied in XTP Multicasting.

4.1 Simulation Run Environment

The performance simulation of XTP Multicasting has been run on the SUN4 processor with the SUNOS4.1 operating system. The standard UNIX debugging tool dbx was used during the upgrade of the XTP simulation program for the inclusion of multicasting.

4.2 The LANSF Tunable Parameters

There are many tunable parameters in the LANSF software. The four parameters which are of interest are mentioned in section 2.4.1 .

4.3 The XTP Tunable Parameters

The seven tunable parameters in this XTP Multicasting simulation are mentioned in section 2.4.2 .

4.4 Measurements

The measurements that are of interest are the throughput and the individual average message delay for the different message sizes under various traffic loads. The strategy for measurements is the same as that applied for XTP Unicast simulation in section 5.4.4 of Chen's thesis [1]. This is summarized in section 2.4.3 of this report.

4.5 Simulation Plan

The simulation plan applied for XTP Multicasting is the same plan as was applied for XTP Unicast simulation in section 5.4.5 of Chen's thesis [1]. Some of these details are mentioned in section 2.4.4 of this report.

The Expected Maximum Throughput = (Message size x Total Bandwidth) / (Total Data Packets Length + Total CNTL Packets Length)

4.6 Simulation Results

The original purpose of the inclusion of XTP V3.3 simulation performance graphs was to compare the XTP V3.5 performance with the XTP V3.3 performance. The XTP V3.3 simulation was run with fixed resend pairs. In the process of progression of this report, the simulation requirements have changed, since it would be appropriate to compare multicasting and unicasting performance at the same revision level XTP V3.5. It was also of interest to run these simulations with no resend pairs.

Hence, in this report, the unicasting and multicasting simulations are based on XTP V3.5 and run **with no resend** pairs, and the simulation results are compared.

4.6.1 XTP V3.3 Unicasting Simulation Results

In Chen [1], the unicasting simulation was based on XTP V3.3 and run with a fixed number of resend pairs (128 bytes). The results obtained are shown in graphs in Figures 1 to 8.

The control packet format was redefined in V3.4, in order to omit transmission of resend pairs where there were no errors. This shortens the length of the body of the control packet by 128 bytes. Throughput is defined as the ratio of the total message size and the total bytes exchanged. For the XTP V3.3 Unicasting, the maximum expected throughput values are tabulated in the second column of Table 17. The maximum expected throughputs with no resend pairs are tabulated in the third column of Table 17, along with the percent improvement shown in the fourth column.

4.6.2 XTP V3.5 Unicasting Simulation Results

Before running the simulations with no resend pairs, the unicasting simulation based on XTP V3.5 was run with fixed number of resend pairs (128 bytes) and compared with the performance of XTP V3.3. The differences found in the performances can be explained in the packet sizes and as per section 2.5.2 .

The results obtained for the throughput and the message delay from the 6 byte message unicasting simulation with no resend pairs are indicated in Tables 1 and 2 respectively. The graphs produced by plotting these results versus the offered load are shown in Figures 1 and 2 respectively. Figures 1 and 2 also show the original graphs respectively for the 6 byte message unicasting simulation, which was based on XTP V3.3.

While the column "No Delay" in Table 1 indicates results when there is no delay due to data copying or checksum calculation, the column "Delay" indicates results when data copying or checksum calculations are carried out.

Similarly, the 128 byte message simulation results are shown in Tables 3 and 4, and Figures 3 and 4 respectively. The 1024 byte message simulation results are shown in Tables 5 and 6, and Figures 5 and 6 respectively. And finally, the 8192 byte message simulation results are shown in Tables 7 and 8, and Figures 7 and 8 respectively.

For the XTP V3.5 Unicasting, the maximum expected throughput values with fixed resend pairs and with no resend pairs, along with the percent improvement are tabulated in Table 18.

4.6.3 XTP V3.5 Multicasting Simulation Results

The multicasting simulation results obtained for the throughput and the message delay are indicated in Tables 9 to 16. The graphs produced by plotting these results versus the offered load are shown in Figures 9 to 16.

For the XTP V3.5 Multicasting, the maximum expected throughput values with fixed resend pairs and with no resend pairs, along with the percent improvement are tabulated in Table 19.

4.6.4 Discussion

For comparing the overall performance of the Unicast XTP V3.3, and Unicast and Multicast XTP V3.5 simulations, the graphs are shown on the same figures respectively. Figures 1 to 8 show the original XTP V3.3 Unicast simulation graphs (with fixed resend pairs) along with the XTP V3.5 Unicasting graphs (with no resend pairs) for 6 byte, 128 byte, 1024 byte and 8192 byte messages respectively.

Since the XTP V3.5 simulation was run with no resend pairs (i.e., 128 bytes are less in the control packet), the total number of bytes in a complete message transmission are reduced. This affects the 6 byte and 128 byte messages where the saturation could not be reached. This required lower message interarrival time. Therefore, three additional testdata files (testdata13 .. testdata15) were added in the case of the 6 byte message, and one additional testdata13 was added in the case of the 128 byte message. Hence, additional three rows in table 1 and 2 for 6 byte and one row in table 3 and 4 were added.

Similarly, Figures 9 to 16 show the original XTP V3.3 Unicast simulation graphs along with the XTP V3.5 Unicasting and Multicasting graphs (with no resend pairs) for 6 byte, 128 byte, 1024 byte and 8192 byte messages respectively.

Note that the scales used in the various graphs are not identical. The 8192 byte messages are fragmented (by LANSF) into five 1460 byte packets and one final packet of 892 bytes. The copy/checksum delay given is the total delay. It is simulated in pieces corresponding to the size of the fragments.

The flat (saturation) parts of the throughput graphs approximates the calculated maximum expected throughput for large message sizes. The throughput graphs in Figures 9, 11, 13, 15 indicate that the flat (saturation) parts of the throughput graphs occur much earlier in the multicast than in the unicast. The multicasting performance is about half of the unicasting performance for small message sizes, and approaches the unicasting performance for large message sizes. The delay graphs in Figures 10, 12, 14 and 16 indicate that the delays are generally higher in multicast than in unicast. It seems clear that the effect of the extra CNTL packet(s) is strong for small message sizes, but not so significant for large message sizes.

The comparison of the unicasting and multicasting maximum expected and actual throughput values with no resend pairs are tabulated in Tables 21 and 22 respectively. The columns 4 in both these tables indicate that the maximum actual multicasting throughput as percent of the actual unicasting throughput approximates the maximum expected multicasting throughput as percentage of maximum expected unicasting throughput.

The number of sample messages being transmitted does not affect the throughput. In XTP Multicast mode, SREQ and EOM are sent in control (CNTL) packets, and not in DATA packets, as it is the case in unicast mode. Therefore, for each SREQ and for each EOM, additional CNTL packets must be sent. For this to happen, the message size must be greater than 1460 bytes.

The individual simulations are subjected to two limits- a limit on the total number of messages sent (10,000 messages) and a limit on the total simulation time. For the unicast case, the message limit is reached first. For the multicast case, the time limit is reached first, and because of the additional CNTL packets transmitted in multicast mode, only about half the test sample messages have been transmitted in the simulation time period. In other words, each CNTL packet uses a message simulation time. For example, in the case of a 6 byte message size in unicast mode, the SREQ/EOM is in the DATA packet. Therefore, only one packet is sent and the total simulation time is used by the total number of test sample messages. In the multicast mode, for every message transmitted, there is at least one control packet sent, which eats up the message simulation time. Therefore, in the total message simulation time, only about half the test sample messages are sent. However, this does not appear to affect the results, as about 5,000 messages are simulated.

Msg length 48 bit (6B) Msg cp delay 16 (time unit) Msg chk delay 1 (time unit)		Effective Throughput with no Resend Pairs (Unicasting XTP V3.5) (TU = Time Unit) (1 sec = 10,000,000TU)			
Msg Inter	Msg Rate	No Delay		Delay	
(Time Unit)	(B/s)	(B/s)	(KB/s)	(B/s)	(KB/s)
22,472	2,730	20,603	20.6	20,603	20.6
11,261	5,328	41,112	41.1	41,115	41.1
7,513	7,986	61,659	61.7	61,292	61.3
5,636	10,644	82,148	82.2	82,148	82.2
4,510	13,302	102,657	102.7	102,657	102.7
3,759	15,960	123,165	123.2	123,165	123.2
3,222	18,618	143,691	143.7	143,691	143.7
2,820	21,276	164,172	164.2	164,172	162.2
2,506	23,934	184,739	184.7	184,739	184.7
2,256	26,582	205,207	205.2	205,207	205.2
2,051	29,250	225,707	225.7	225,708	225.7
1,880	31,908	246,117	246.1	246.116	246.1
1,735	34,566	266,117	266.2	266,232	266.2
1,611	37,244	269,407	269.4	269,408	269.4
1,504	39894	269,570	269.6	269,569	269.6
1,410	42,553	269,698	269.7	269,698	269.7

Table 1: Throughput vs Offered load for 6 byte messages with no Resend Pairs in Unicasting XTP V3.5

Msg length 48 bit (6B) Msg cp delay 16 (time unit) Msg chk delay 1 (time unit)		Message Delay with no Resend Pairs (Unicasting XTP V3.5) (TU = Time Unit) (1 sec = 10,000,000TU)					
Msg Inter (TU)	Msg Rate (B/s)	No Delay			Delay		
		Min Delay (TU)	Msg Delay		Min Delay (TU)	Msg Delay	
			(TU)	(ms)		(TU)	(ms)
22,472	2,730	1607	213	0.02	1622	215	0.02
11,261	5,328	1607	208	0.02	1622	210	0.02
7,513	7,986	1607	208	0.02	1622	211	0.02
5,636	10,644	1607	208	0.02	1622	210	0.02
4,510	13,302	1607	208	0.02	1622	210	0.02
3,759	15,960	1607	208	0.02	1622	210	0.02
3,222	18,618	1607	208	0.02	1622	210	0.02
2,820	21,276	1607	208	0.02	1622	210	0.02
2,506	23,934	1607	208	0.02	1622	210	0.02
2,256	26,582	1607	208	0.02	1622	210	0.02
2,051	29,250	1607	208	0.02	1622	210	0.02
1,880	31,908	1607	208	0.02	1622	210	0.02
1,735	34,566	1607	209	0.02	1622	211	0.02
1,611	37,244	1609	222	0.02	1623	228	0.02
1,504	39,894	1609	223	0.02	1624	239	0.02
1,410	42,553	1609	254	0.03	1624	223	0.02

Table 2: Delay vs Offered load for 6 byte messages with no Resend Pairs in Unicasting XTP V3.5

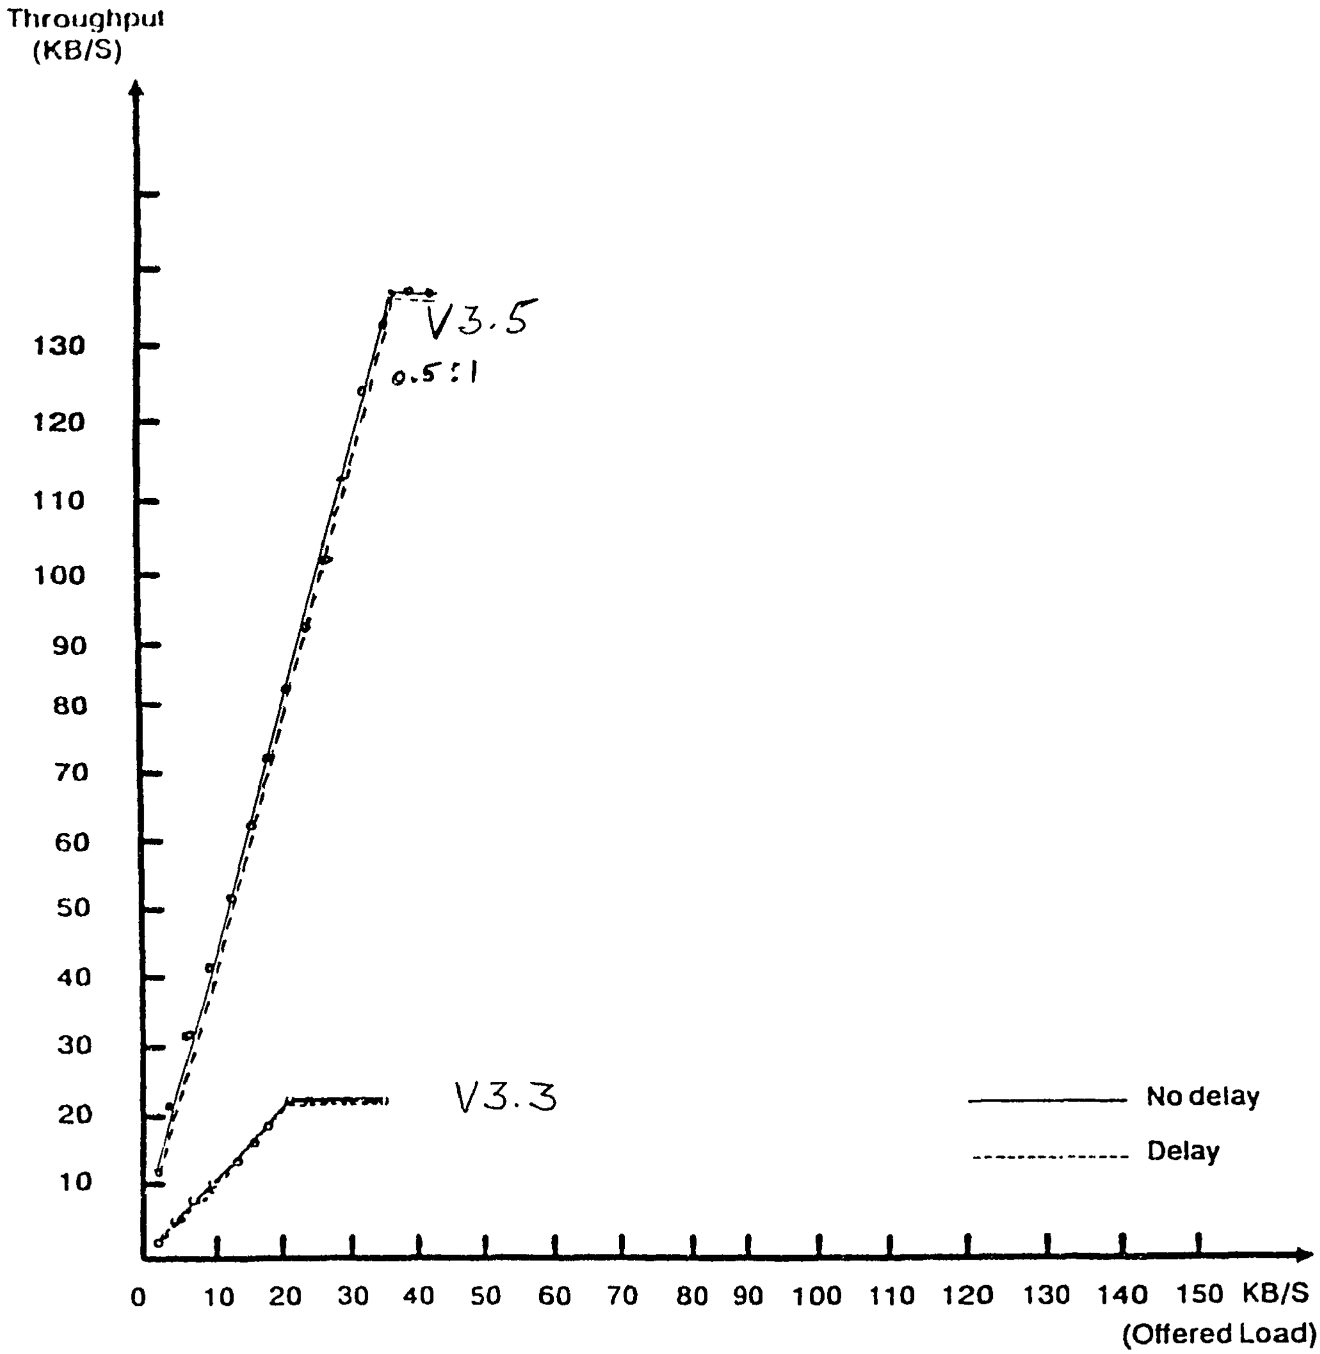

Figure 1: Throughput vs offered load for 6 byte messages with no Resend Pairs in Unicasting XTP V3.5

V3.3 - Unicasting with resend pairs

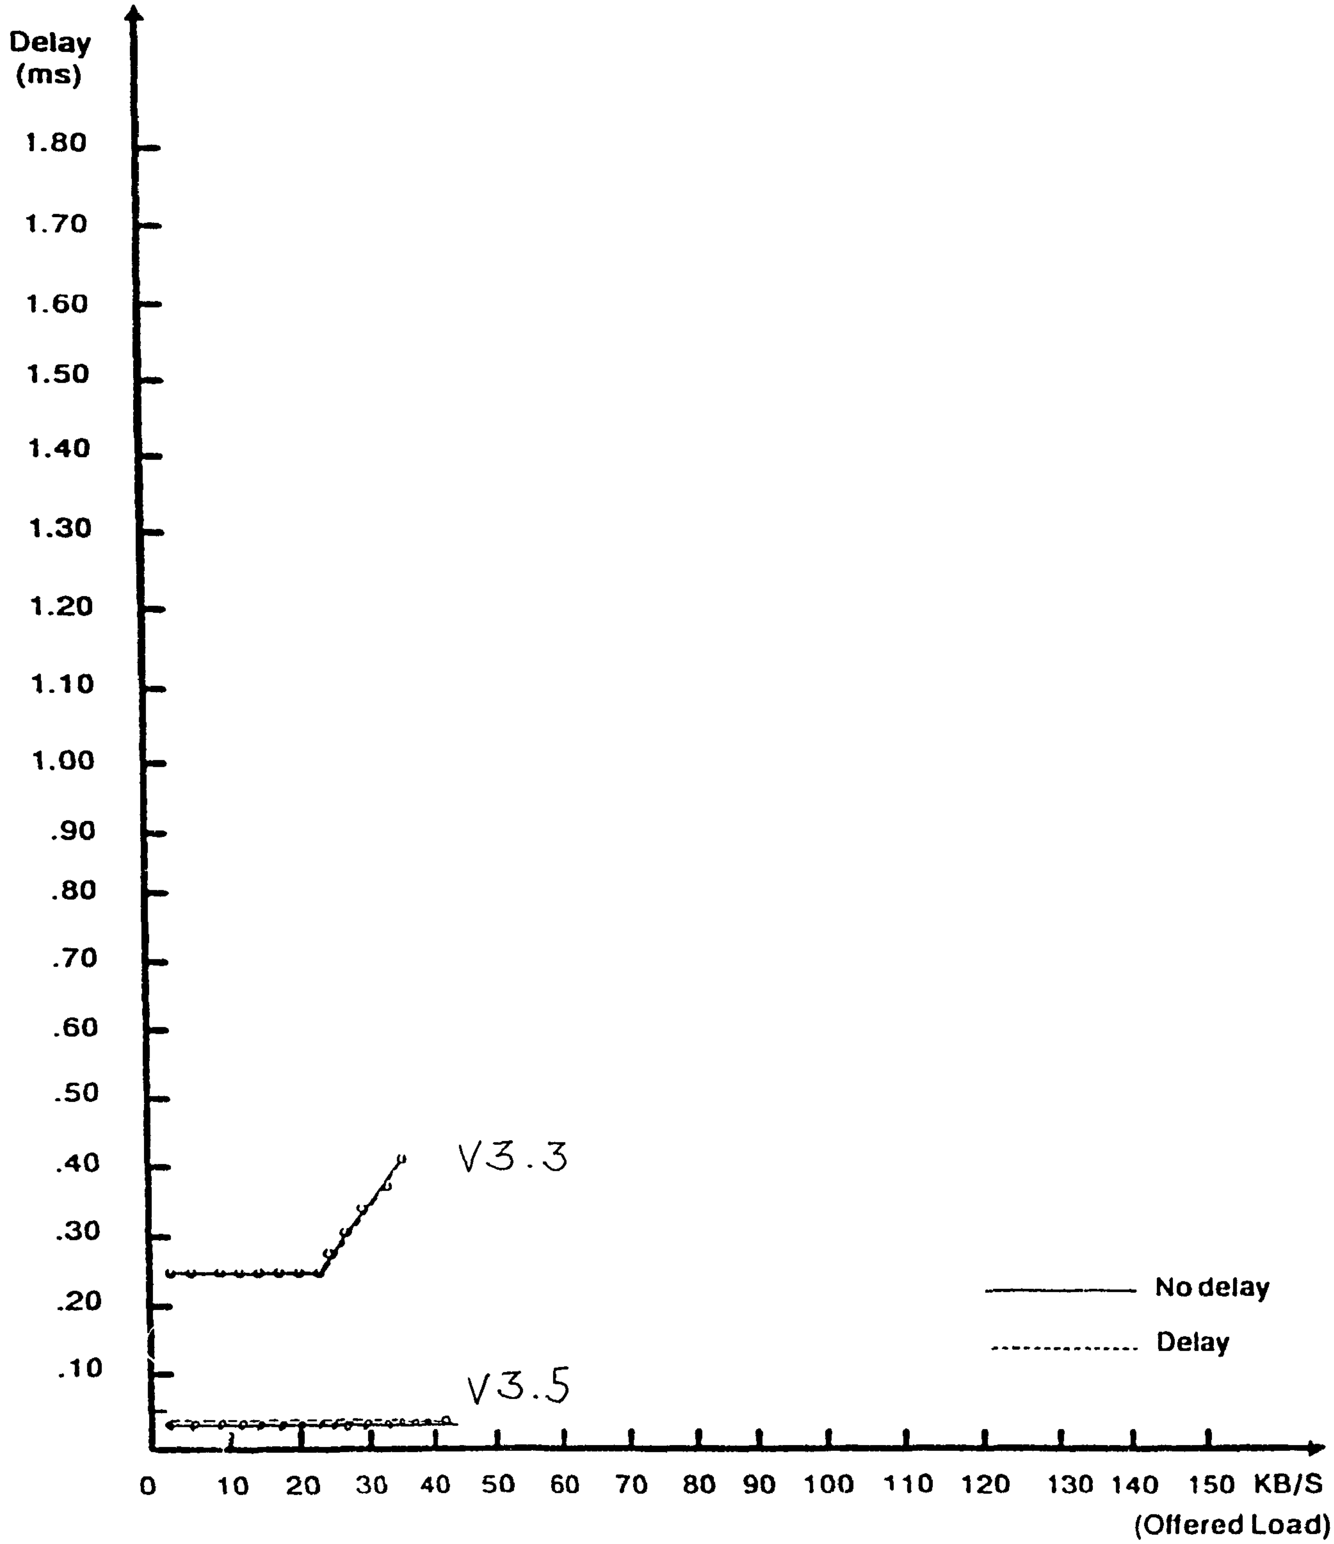

Figure 2: Delay vs offered load for
6 byte messages with no Resend Pairs
in Unicasting XTP V3.5

V3.3 - Unicasting with resend pairs

Msg length 1024 bit (128B) Msg cp delay 336 (time unit) Msg chk delay 1 (time unit)		Effective Throughput with no Resend Pairs (Unicasting XTP V3.5) (TU = Time Unit) (1 sec = 10,000,000TU)			
Msg Inter	Msg Rate	No Delay		Delay	
(Time Unit)	(B/s)	(B/s)	(KB/s)	(B/s)	(KB/s)
32,258	39,680	52,420	52.4	52,420	52.4
16,077	79,616	105,178	105.2	105,178	105.2
10,741	119,168	157,428	157.4	157,427	157.4
8,064	158,720	209,686	209.7	209,684	209.7
6,455	198,272	261,950	262.0	261,946	262.0
5,382	237,824	314,170	314.2	314,165	314.2
4,615	277,776	366,378	366.4	366,369	366.4
4,038	316,928	418,723	418.7	418,713	418.7
3,591	356,352	470,834	470.8	470,820	470.8
3,232	396,032	523,112	523,1	522,843	522.8
2,937	435,712	475,328	575.3	570,417	570.4
2,692	475,392	623,929	624.0	571,531	571.5
2,45	515,072	625,853	625.9	571,392	571.4
2,309	554,353	626,224	626.2	571,266	571.3

Table 3: Throughput vs Offered load for
128 byte messages with no Resend Pairs
in Unicasting XTP V3.5

Msg length 1024 bit (128B) Msg cp delay 336 (time unit) Msg chk delay 1 (time unit)		Message Delay with no Resend Pairs (Unicasting XTP V3.5) (TU = Time Unit) (1 sec = 10,000,000TU)					
Msg Inter (TU)	Msg Rate (B/s)	No Delay			Delay		
		Min Delay (TU)	Msg Delay (TU)	Msg Delay (ms)	Min Delay (TU)	Msg Delay (TU)	Msg Delay (ms)
32,258	39,680	2583	1955	0.20	2824	2138	0.21
16,077	79,616	2583	1955	0.20	2824	2138	0.21
10,741	119,168	2583	1955	0.20	2824	2138	0.21
8,064	158,720	2583	1955	0.20	2824	2138	0.21
6,455	198,272	2583	1955	0.20	2824	2138	0.21
5,382	237,824	2583	1955	0.20	2824	2138	0.21
4,615	277,776	2583	1958	0.20	2824	2141	0.21
4,038	316,928	2583	1955	0.20	2824	2138	0.21
3,591	356,352	2583	1955	0.20	2824	2138	0.21
3,232	396,032	2583	1956	0.20	2824	2138	0.21
2,937	435,712	2583	9156	0.20	2824	2157	0.22
2,692	475,392	2583	1968	0.20	2824	2344	0.24
2,485	515,072	2583	2126	0.21	2824	2546	0.26
2,309	554,353	2585	2288	0.23	2824	2740	0.27

Table 4: Delay vs Offered load for 128 byte messages with no Resend Pairs in Unicasting XTP V3.5

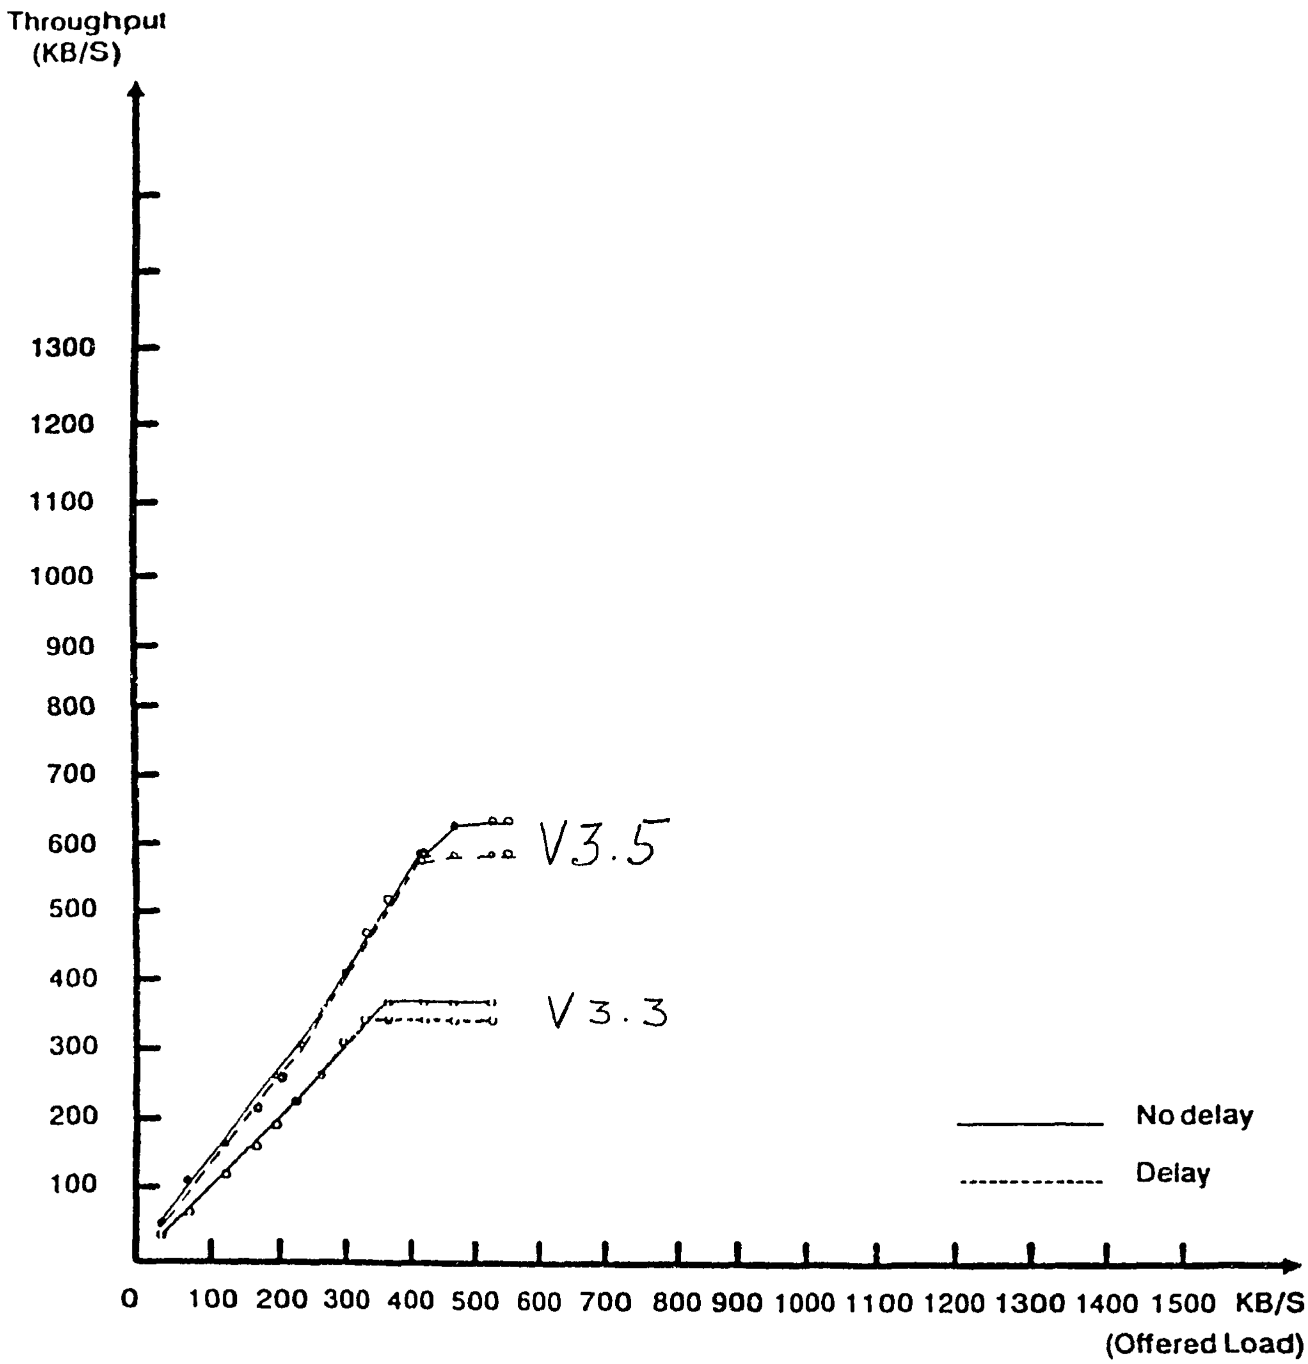

Figure 3: Throughput vs offered load for 128 byte messages with no Resend Pairs in Unicasting XTP V3.5

V3.3 - Unicasting with resend pairs

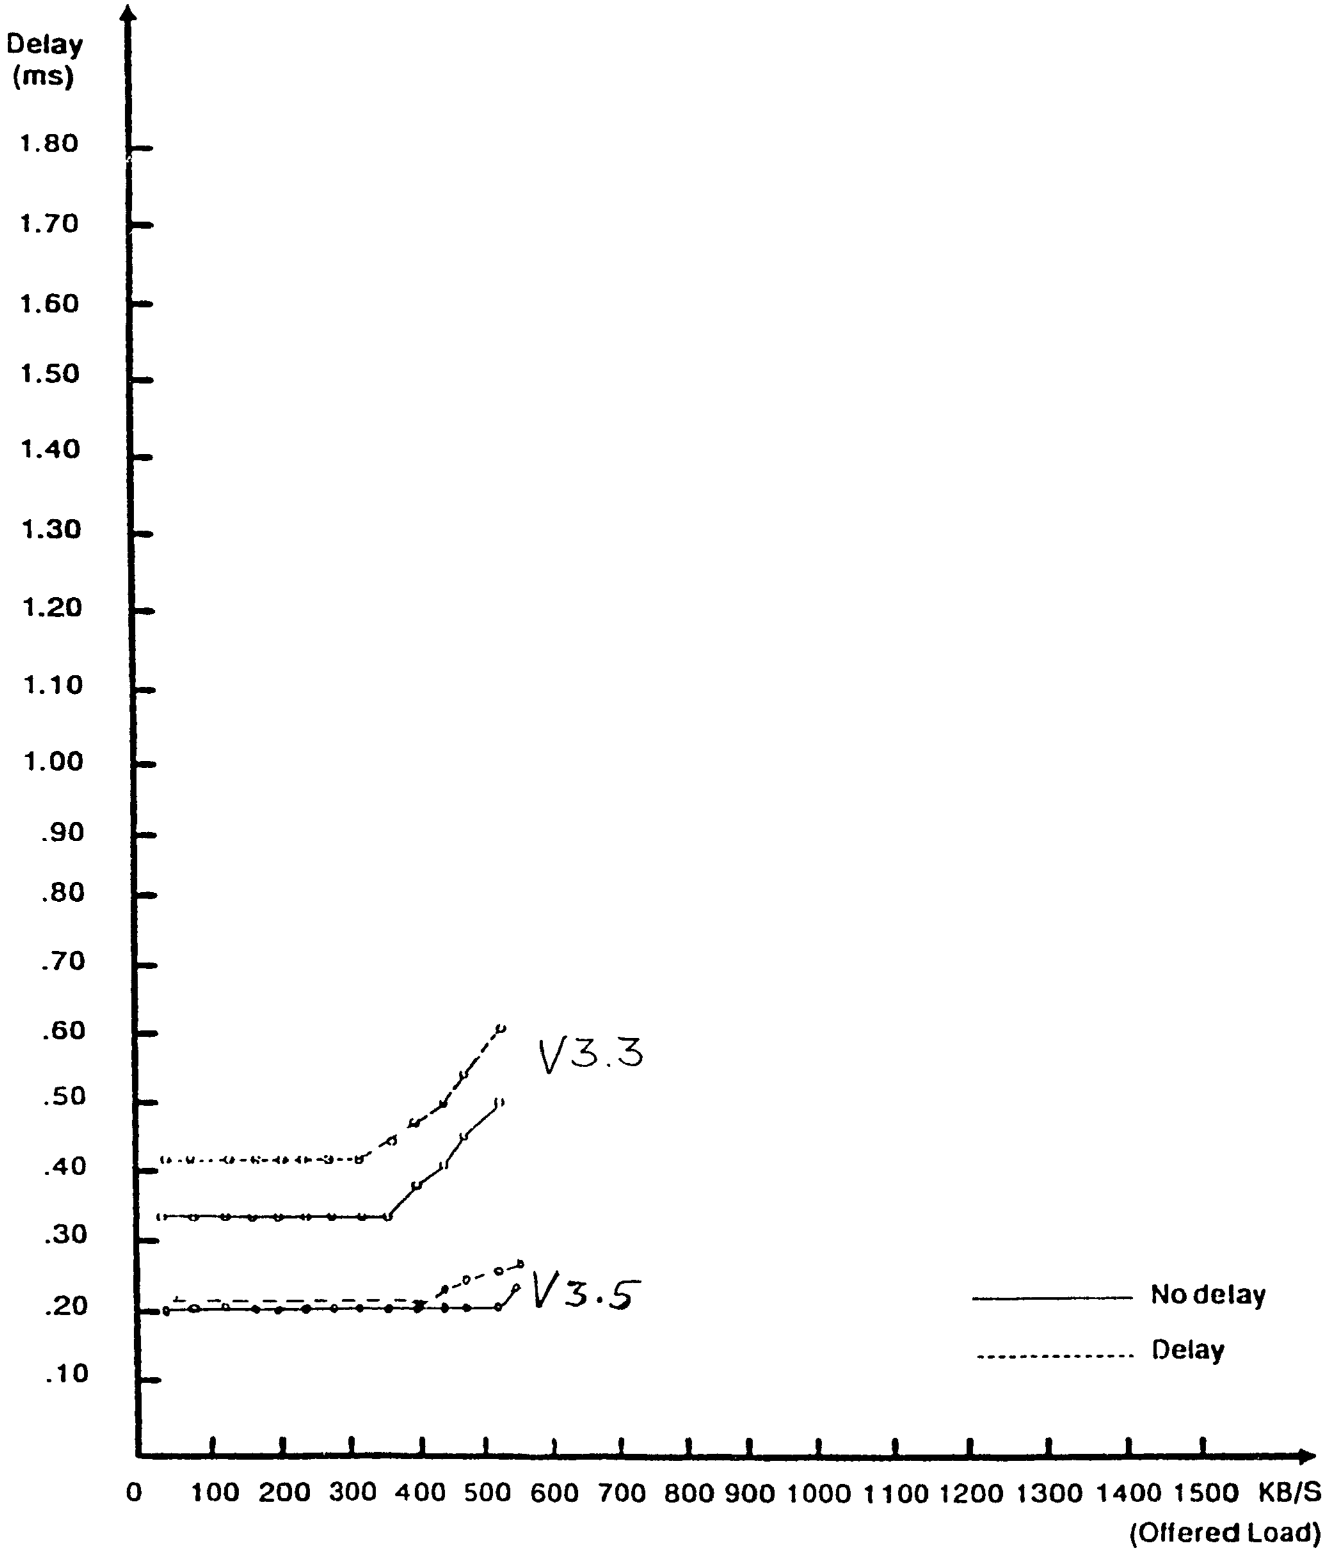

Figure 4: Delay vs offered load for
128 byte messages with no Resend Pairs
in Unicasting XTP V3.5

V3.3 - Unicasting with resend pairs

Msg length 8192 bit (1KB) Msg cp delay 2,694 (time unit) Msg chk delay 1 (time unit)		Effective Throughput with no Resend Pairs (Unicasting XTP V3.5) (TU = Time Unit) (1 sec = 10,000,000TU)			
Msg Inter	Msg Rate	No Delay		Delay	
(Time Unit)	(B/s)	(B/s)	(KB/s)	(B/s)	(KB/s)
104,166	98,324	102,811	102.8	102,811	102.8
51,813	197,632	206,691	206.7	207,355	206.4
34,602	295,936	309,495	309.5	309,487	309.5
25,974	394,240	412,297	412.3	412,283	412.3
20,790	492,544	515,095	515.1	515,069	515.1
17,331	590,848	652,300	652.3	617,850	617.9
14,85[illegible]	689,152	720,717	720.7	720,569	720.6
13,003	787,456	823,512	823.5	821,316	823.3
11,645	875,520	919,445	919.5	824,897	824.9
10,405	984,064	1,028,102	1028.1	821,502	821.5
9,460	1,082,368	1,078,331	1078.3	817,402	817.4
8,673	1,184,748	1,079,424	1079.4	812,186	812.2
8,006	1,278,976	1,080,019	1080.0	807,431	807.4

Table 5: Throughput vs Offered load for 1 Kbyte messages with no Resend Pairs in Unicasting XTP V3.5

Msg length 8192 bit (1KB) Msg cp delay 21,559 (TU) Msg chk delay 1 (time unit)		Message Delay with no Resend Pairs (Unicasting XTP V3.5) (TU = Time Unit) (1 sec = 10,000,000TU)					
Msg Inter (TU)	Msg Rate (B/s)	No Delay			Delay		
		Min Delay (TU)	Msg Delay		Min Delay (TU)	Msg Delay	
			(TU)	(ms)		(TU)	(ms)
104,166	98,324	9750	9325	0.93	12349	11810	1.18
51,813	197,632	9750	9323	0.93	12349	11770	1.18
34,604	295,936	9750	9323	0.93	12349	11808	1.18
25,974	394,240	9750	9323	0.93	12349	11808	1.18
20,730	492,544	9750	9323	0.93	12349	11808	1.18
17,331	590,848	9750	8832	0.93	12349	11809	1.18
14,858	689,152	9750	9323	0.93	12349	11811	1.18
13,003	787,456	9750	9323	0.93	12349	11840	1.18
11,645	875,520	9750	9284	0.93	12349	13107	1.31
10,405	984,064	9750	9332	0.93	12349	14793	1.48
9,460	1082,368	9750	9789	0.93	12349	16352	1.64
8,673	1184,748	9750	10704	0.93	12349	18014	1.80
8,006	1278,976	9750	11584	0.93	12349	19561	1.96

Table 6: Delay vs Offered load for 1 Kbyte messages with no Resend Pairs in Unicasting XTP V3.5

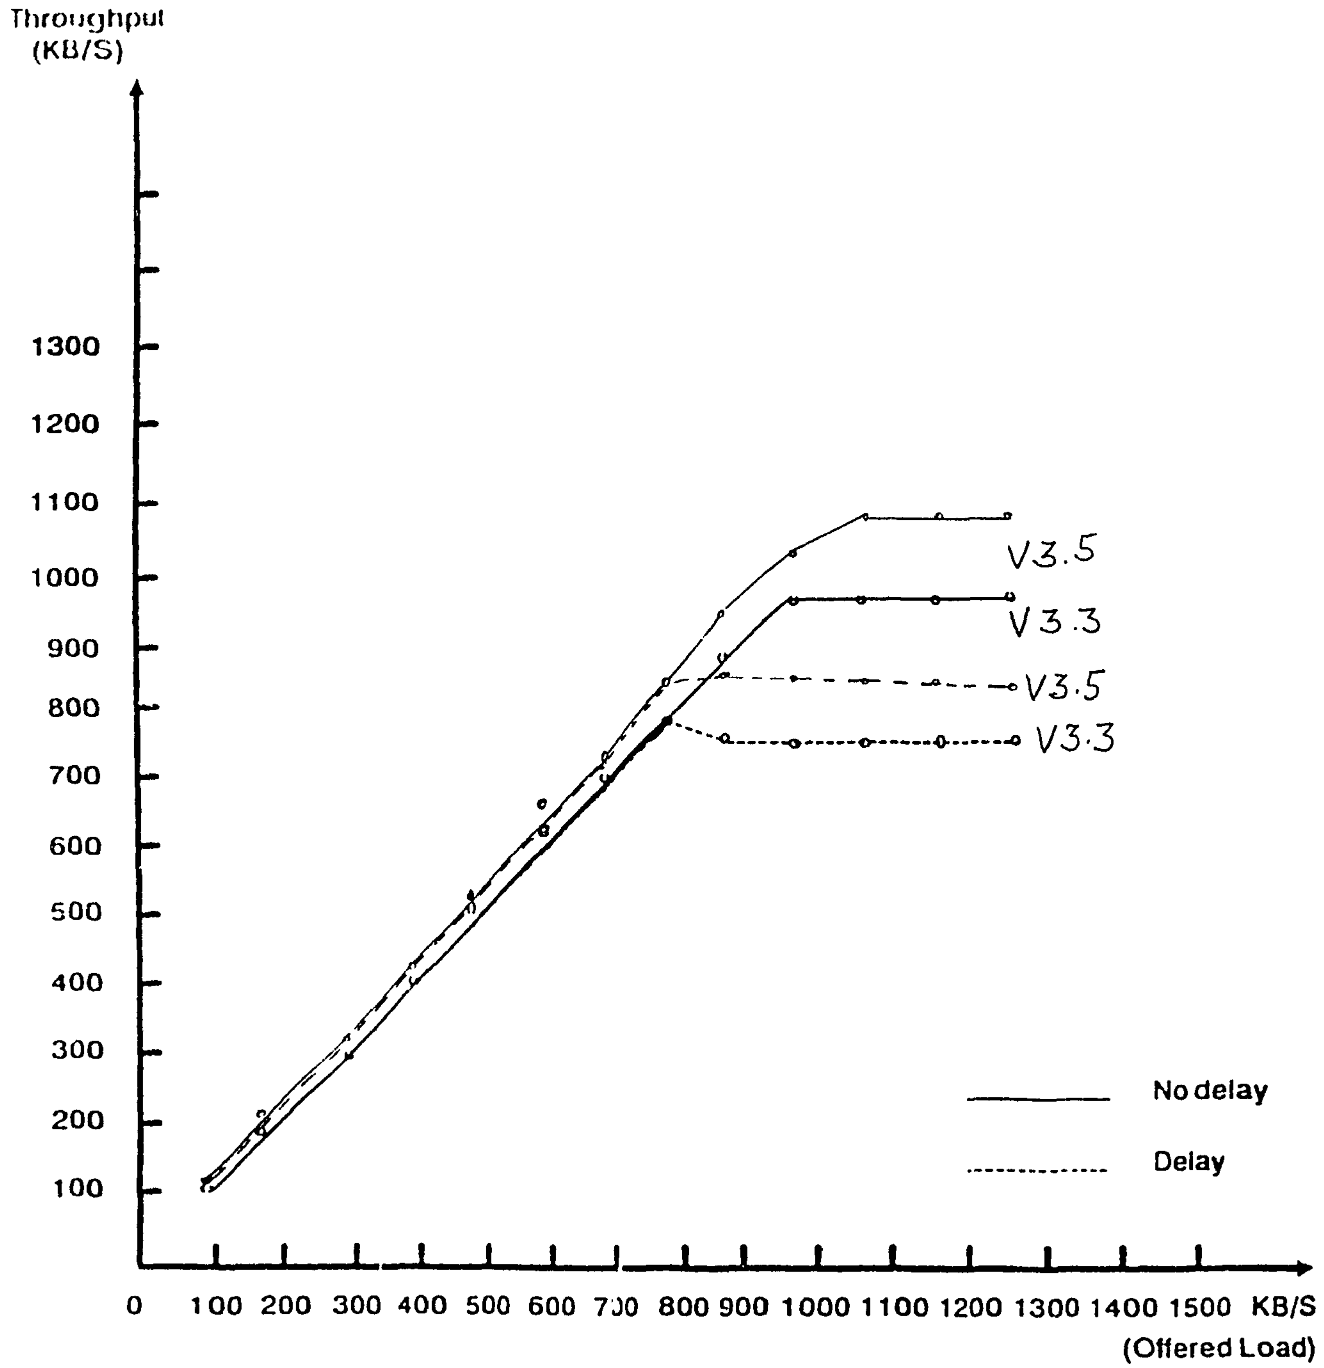

Figure 5: Throughput vs offered load for 1 Kbyte messages with no Resend Pairs in Unicasting XTP V3.5

V3.3 - Unicasting with resend pairs

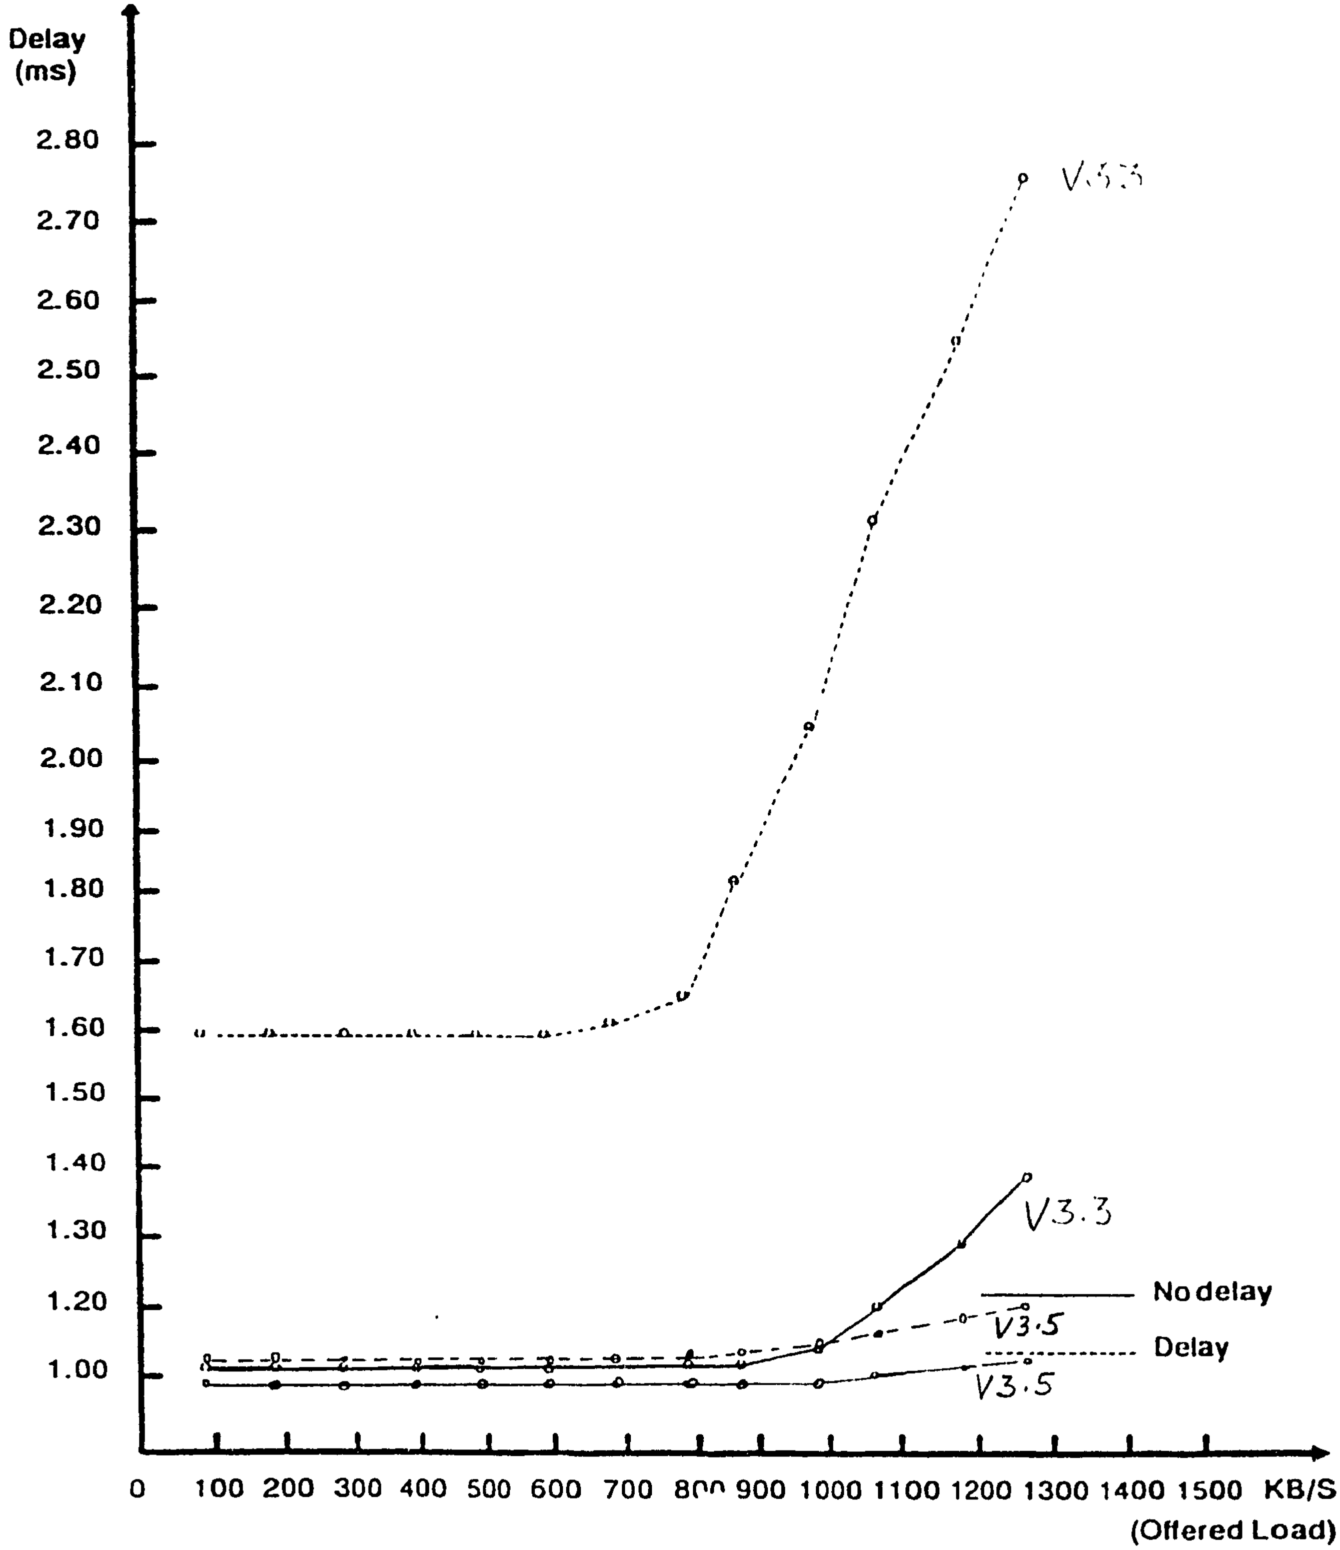

Figure 6: Delay vs offered load for
1 Kbyte messages with no Resend Pairs
in Unicasting XTP V3.5

V3.3 - Unicasting with resend pairs

Msg length 65,536 bit (8KB) Msg cp delay 21,559 (time unit) Msg chk delay 1 (time unit)		Effective Throughput with no Resend Pairs (Unicasting XTP V3.5) (TU = Time Unit) (1 sec = 10,000,000TU)			
Msg Inter	Msg Rate	No Delay		Delay	
(Time Unit)	(B/s)	(B/s)	(K/s)	(B/s)	(KB/s)
714,285	114,688	117,085	117.1	117,048	117.1
357,142	229,376	236,935	236.9	236,851	236.9
238,095	344,064	355,373	355.4	355,449	355.5
178,571	458,752	473,852	473.9	473,341	473.3
142,857	573,440	593,795	593.8	592,225	592.2
119,047	688,128	710,234	710.2	712,772	712.8
104,166	786,432	814,293	814.3	814,192	814.2
89,285	917,502	949,521	949.5	949,097	949.1
79,363	1,032,192	1067,713	1067.7	1052,162	1052.2
71,428	1,146,880	1176,225	1176.2	1049,850	1044.9
64,935	1,261,568	1178,192	1178.2	1047,378	1047.4
59,524	1,376,256	1179,297	1179.3	1043,155	1043.2
54,945	1,490,944	1180,449	1180.5	1039,283	1039.3

Table 7: Throughput vs Offered load for 8 Kbyte messages with no Resend Pairs in Unicasting XTP V3.5

Msg length 65,536 bit (8KB) Msg cp delay 21,559 (TU) Msg chk delay 1 (time unit)		Message Delay with no Resend Pairs (Unicasting XTP V3.5) (TU = Time Unit) (1 sec = 10,000,000TU)					
Msg Inter (TU)	Msg Rate (B/s)	No Delay			Delay		
		Min Delay (TU)	Msg Delay		Min Delay (TU)	Msg Delay	
			(TU)	(ms)		(TU)	(m)
714,285	114,688	8,788	8,608	0.86	10.947	10,726	1.07
357,142	229,376	8,788	8,508	0.85	10,947	10,602	1.06
238,095	344,064	8,787	8,507	0.85	10,947	10,596	1.06
178,571	458,752	8,787	8,507	0.85	10,947	10,610	1.06
142,857	573,440	8,787	8,486	0.85	10,947	10,600	1.06
119,047	688,128	8,787	8,514	0.85	10,947	10,569	1.06
104,166	786,432	8,787	8,486	0.85	10,947	10,574	1.06
89,285	917,502	8,787	8,491	0.85	10,947	10,583	1.06
79,363	1032,192	8,787	8,495	0.85	10,947	10,739	1.07
71,428	1146,880	8,787	8,568	0.86	10,947	11,959	1.20
64,935	1261,568	8,787	9,409	0.94	10,947	13,186	1.32
59,524	1376,256	8,787	10,255	1.03	10,947	14,443	1.44
54,945	1490,944	8,787	11,098	1.11	10,947	15705	1.57

Table 8: Delay vs Offered load for
8 Kbyte messages with no Resend Pairs
in Unicasting XTP V3.5

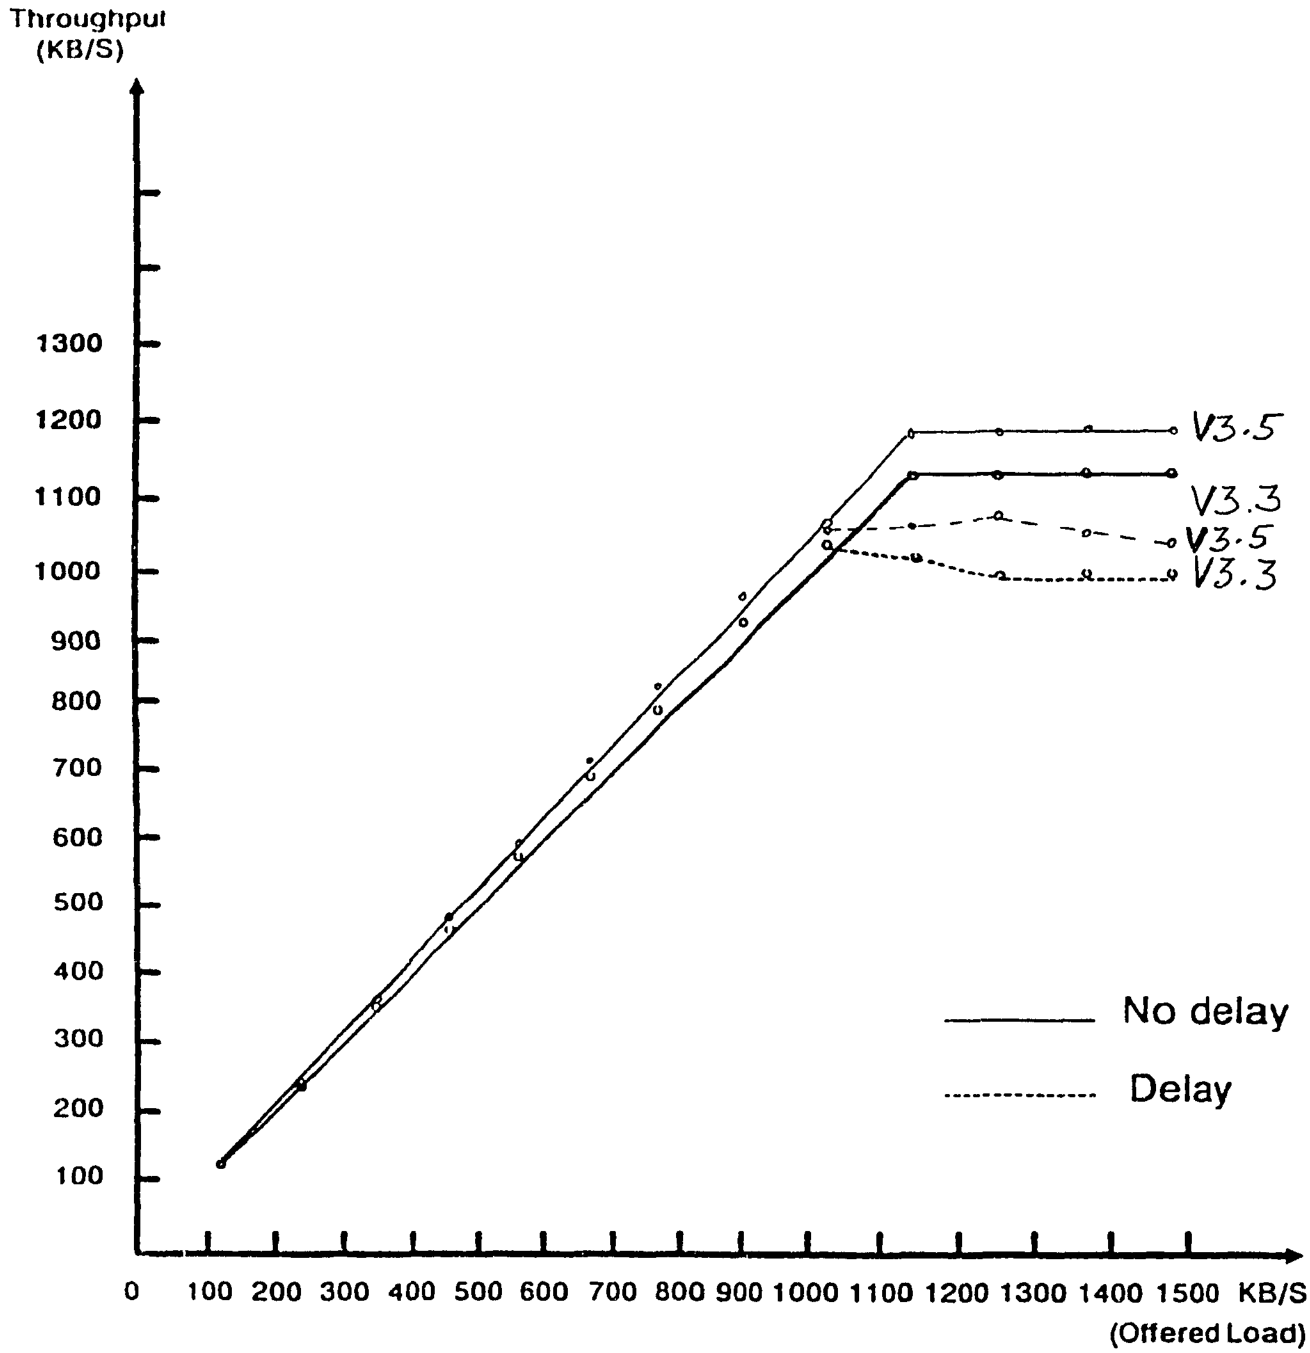

Figure 7: Throughput vs offered load for 8 Kbyte messages with no Resend Pairs in Unicasting XTP V3.5

V3.3 - Unicasting with resend pairs

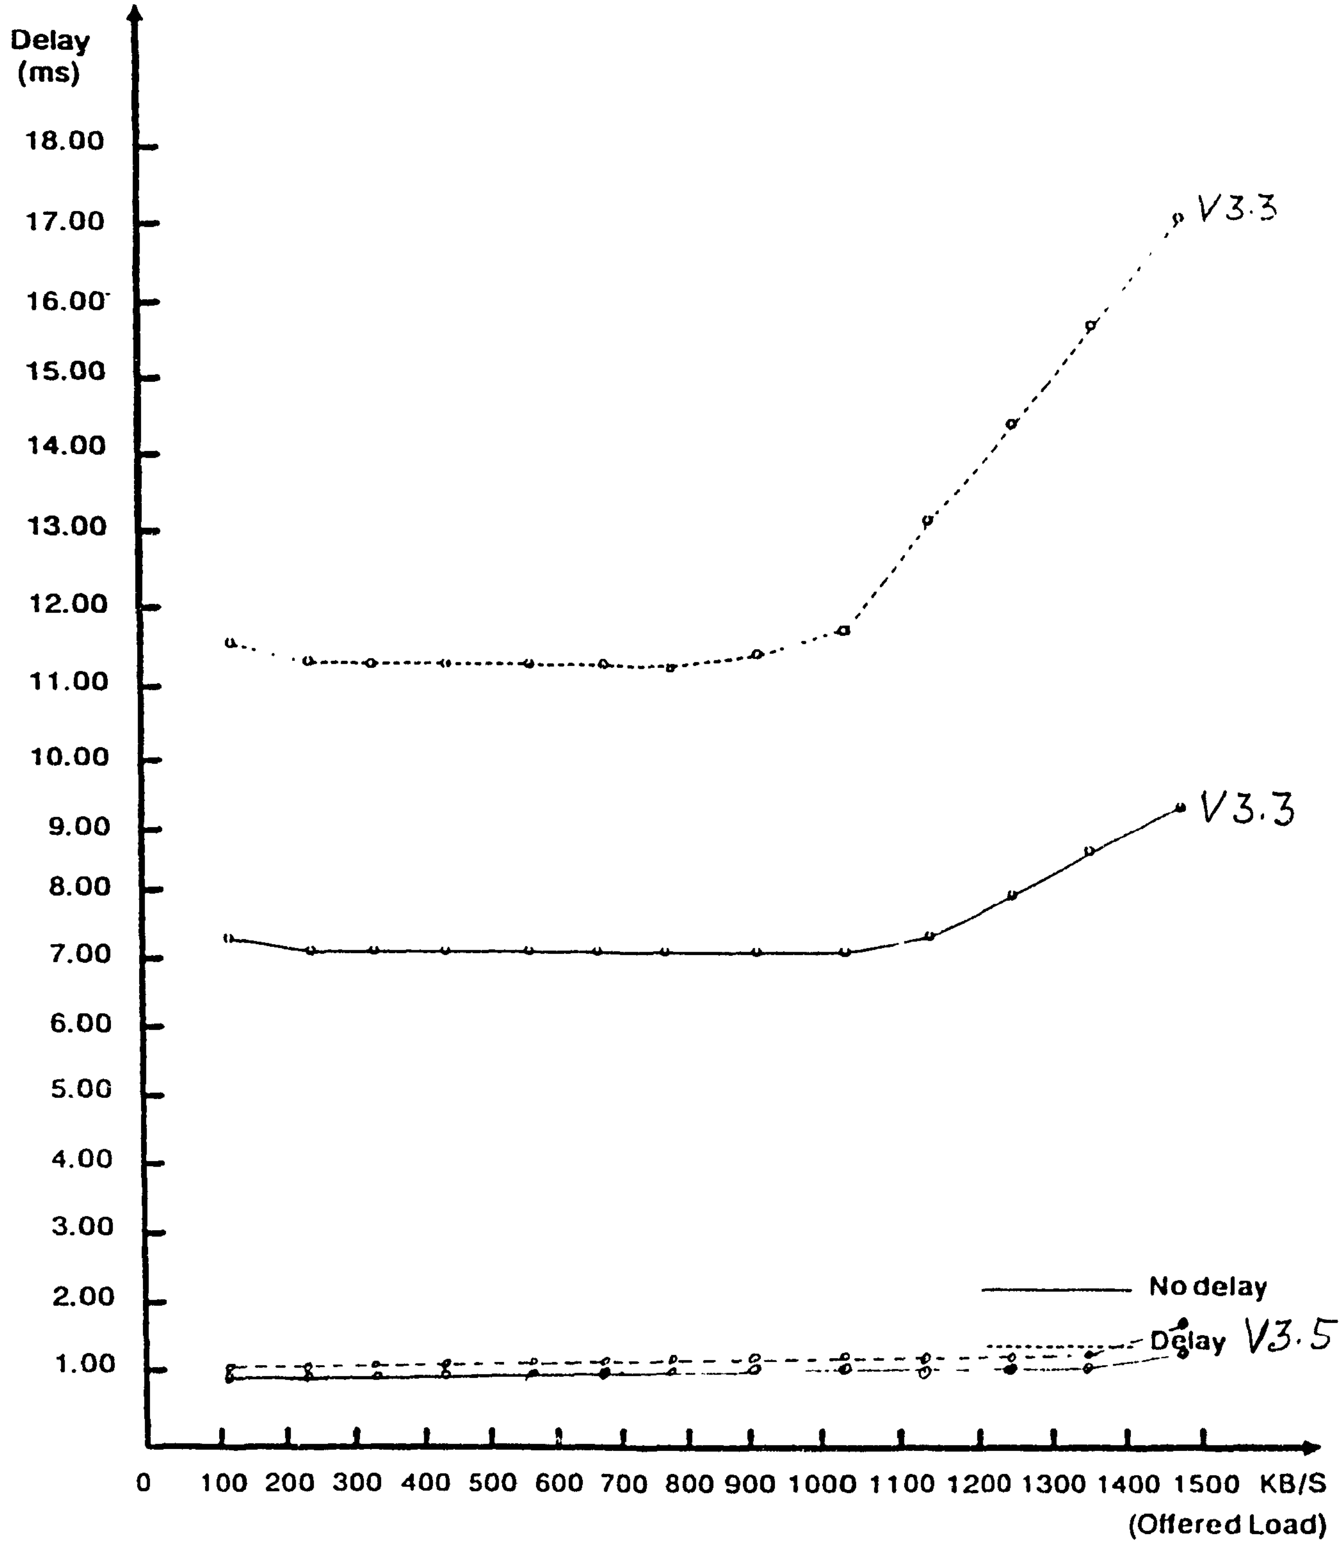

Figure 8: Delay vs offered load for
8 Kbyte messages with no Resend Pairs
in Unicasting XTP V3.5

V3.3 - Unicasting with resend pairs

Msg length 48 bit (6B) Msg cp delay 16 (time unit) Msg chk delay 1 (time unit)		Effective Throughput with no Resend Pairs (Multicasting XTP V3.5) (TU = Time Unit) (1 sec = 10,000,000TU)			
Msg Inter	Msg Rate	No Delay		Delay	
(Time Unit)	(B/s)	(B/s)	(KB/s)	(B/s)	(KB/s)
22,472	2,730	20,556	20.6	20,556	20.6
11,261	5,328	41,019	41.0	41,019	41.0
7,513	7,986	61,480	61.5	61,480	61.5
5,636	10,644	81,951	82.0	81,951	82.0
4,510	13,302	102,402	102.4	102,403	102.4
3,759	15,960	122,852	122.9	122,852	122.9
3,222	18,618	143,307	143.3	143,307	143.3
2,820	21,276	163,469	163.5	163,469	163.5
2,506	23,934	169,272	169.3	169,273	164.3
2,256	26,582	169,293	169.3	169,293	169.3
2,051	29,250	169,311	169.3	169,311	169.3
1,880	31,908	169,325	169.3	169,325	169.3
1,735	34,566	169.336	169.3	169,335	169.3

Table 9: Throughput vs Offered load for 6 byte messages with no Resend Pairs in Multicasting XTP V3.5

Msg length 48 bit (6B) Msg cp delay 16 (time unit) Msg chk delay 1 (time unit)		Message Delay with no Resend Pairs (Multicasting XTP V3.5) (TU = Time Unit) (1 sec = 10,000,000TU)					
Msg Inter (TU)	Msg Rate (B/s)	No Delay			Delay		
		Min Delay (TU)	Msg Delay		Min Delay (TU)	Msg Delay	
			(TU)	(ms)		(TU)	(ms)
22,472	2,730	2615	347.3	0.04	2630	349.3	0.04
11,261	5,328	2615	339.7	0.03	2630	341.6	0.03
7,513	7,986	2615	339.7	0.03	2630	341.6	0.03
5,636	10,644	2435	316.3	0.03	2450	318.2	0.03
4,510	13,302	2615	339.7	0.03	2630	341.6	0.03
3,759	15,960	2615	339.7	0.03	2630	341.7	0.03
3,222	18,618	2615	339.7	0.03	2631	341.8	0.03
2,820	21,276	2616	340.5	0.03	2631	342.4	0.03
2,506	23,934	2649	374.6	0.04	2656	375.5	0.04
2,256	26,582	2650	416.1	0.04	2655	416.9	0.04
2,051	29,250	2648	457.5	0.05	2654	458.5	0.05
1,880	31,908	2618	493.3	0.05	2632	496.0	0.05
1,735	34,566	2679	546.9	0.06	2656	542.2	0.05

Table 10: Delay vs Offered load for
6 byte messages with no Resend Pairs
in Multicasting XTP V3.5

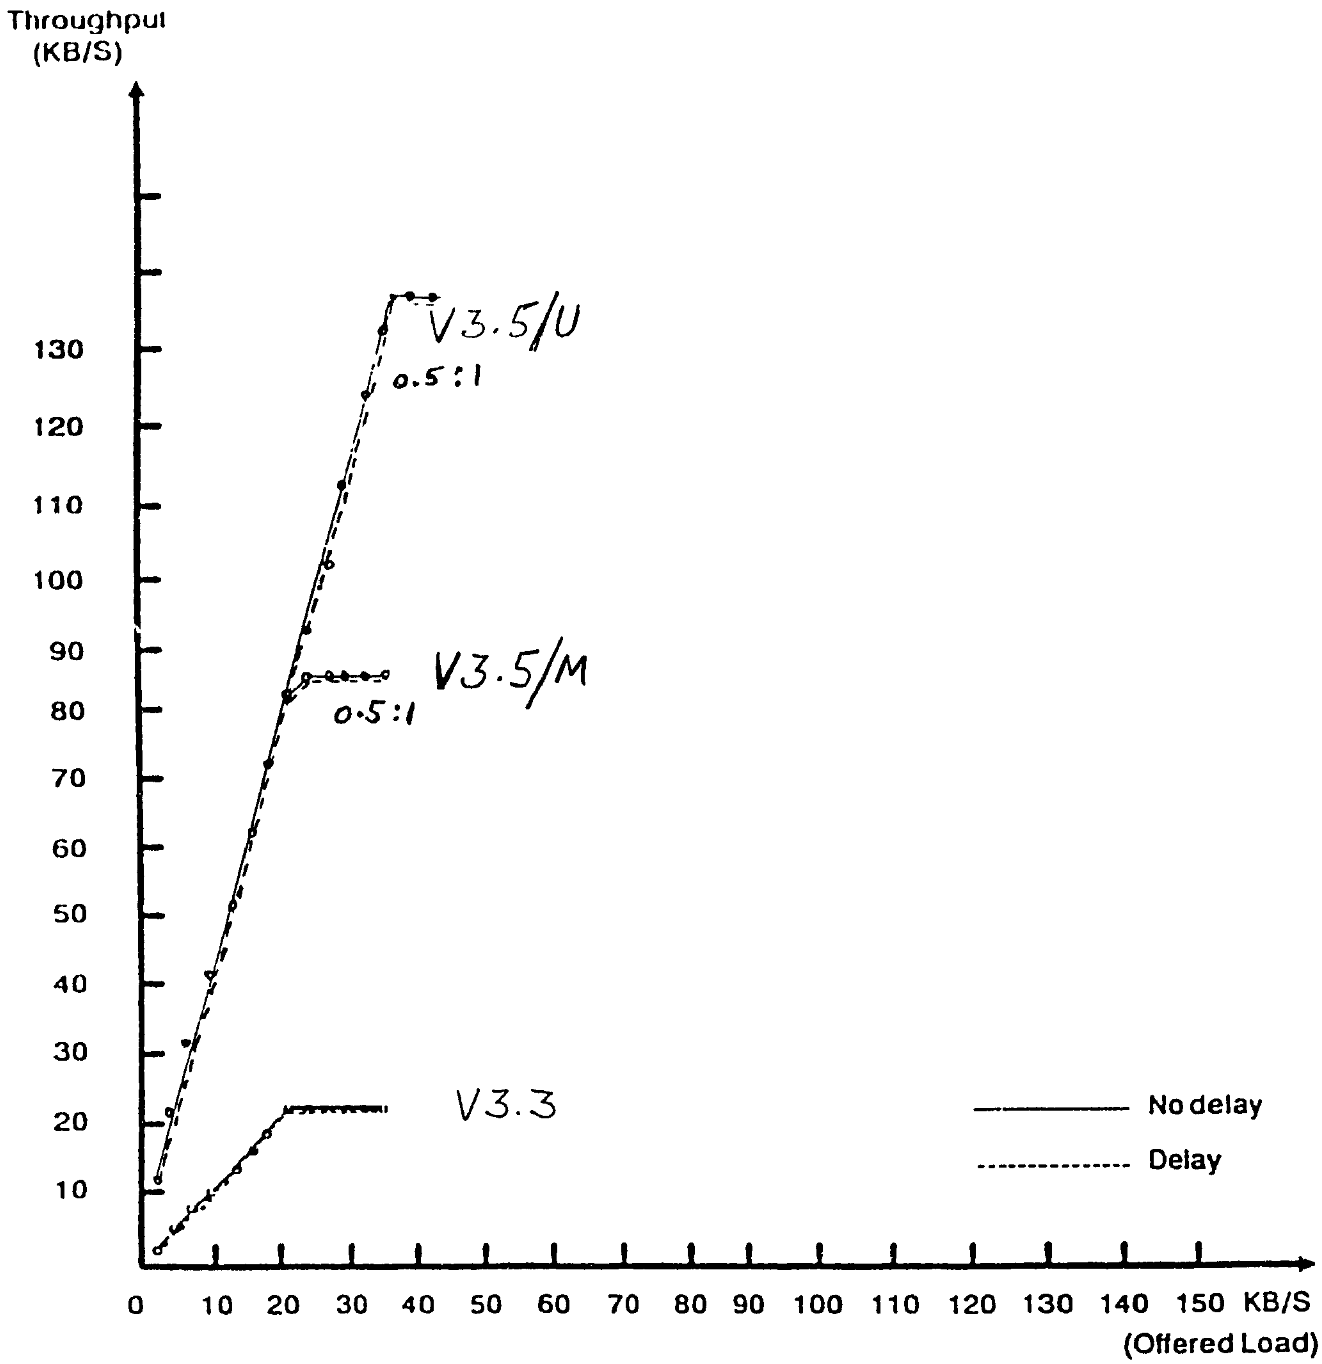

Figure 9: Throughput vs offered load for 6 byte messages with no Resend Pairs in Multicasting XTP V3.5

V3.5/M - Multicasting with no resend pairs
V3.5/U - Unicasting with no resend pairs
3.3 - Unicasting with resend pairs

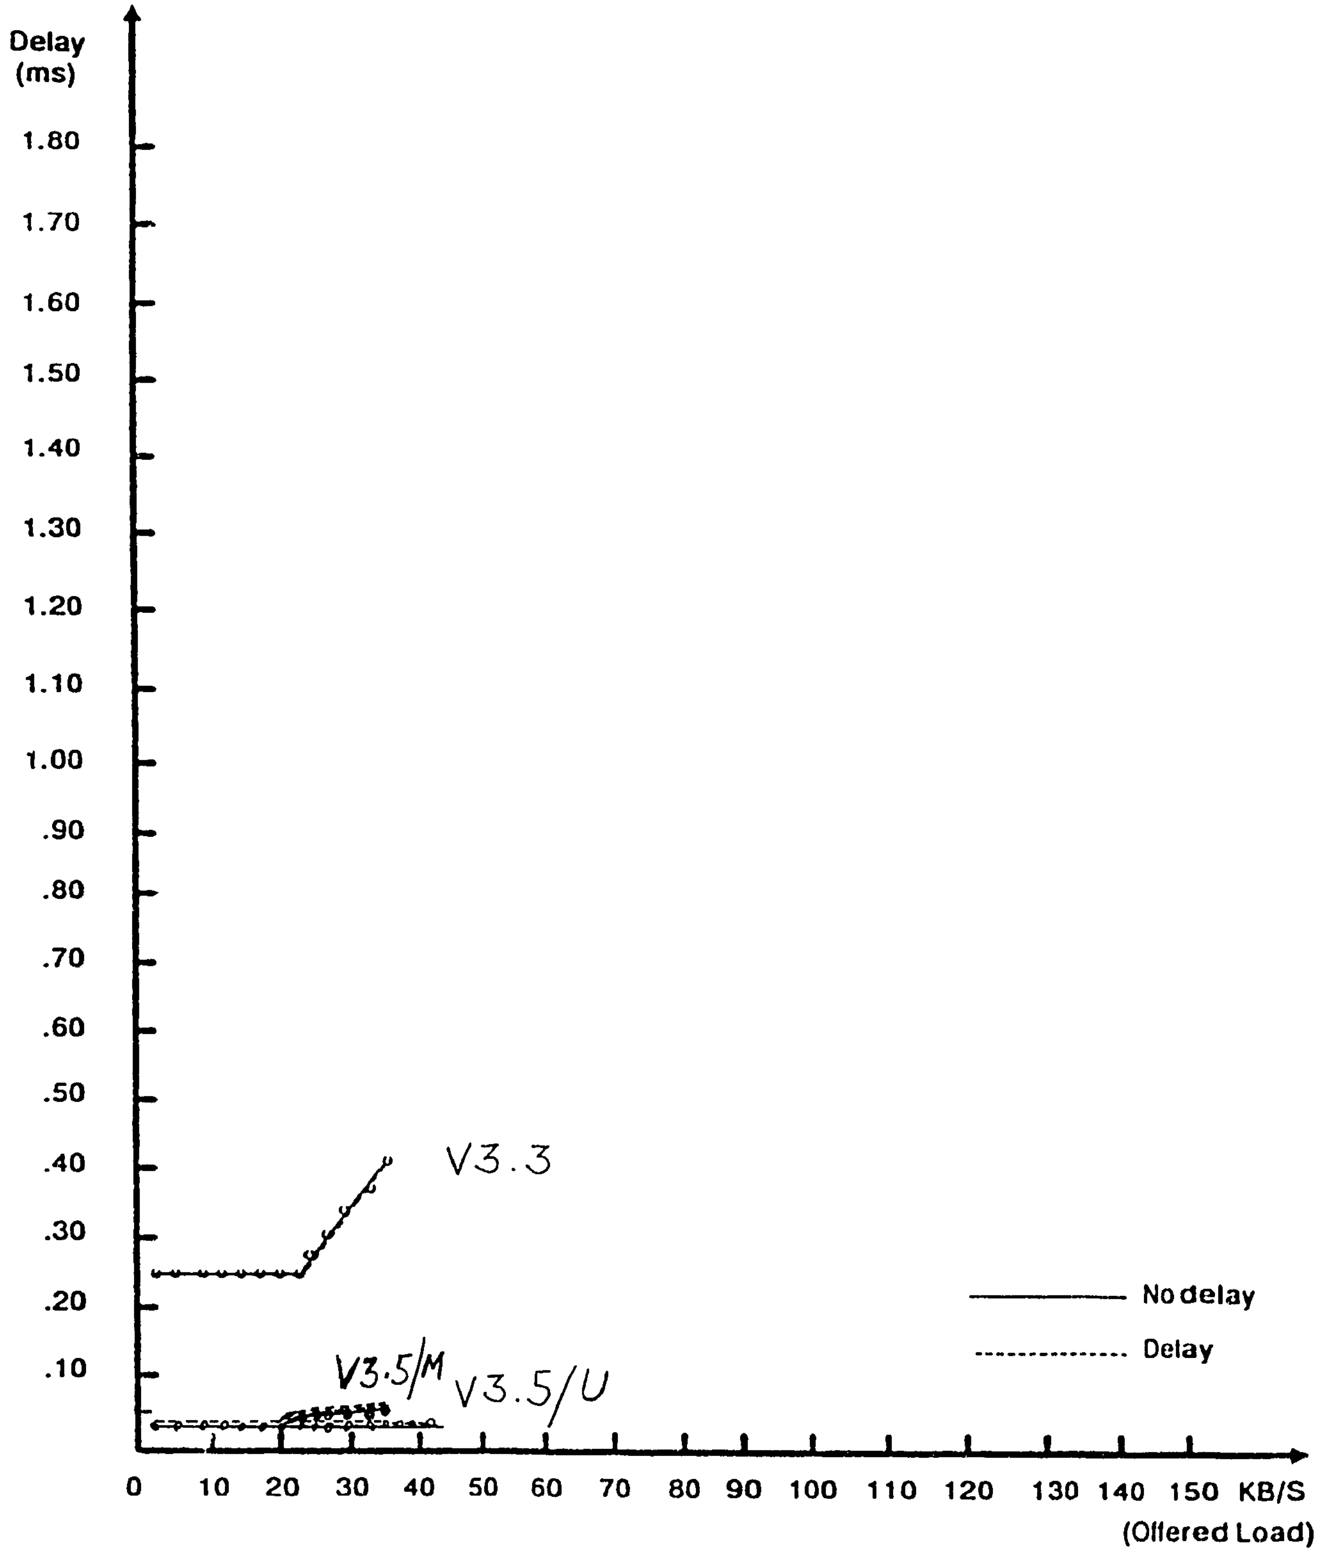

Figure 10: Delay vs offered load for
6 byte messages with no Resend Pairs
in Multicasting XTP V3.5

V3.3/M - Multicasting with no resend pairs
V3.5/U - Unicasting with no resend pairs
V3.3 - Unicasting with resend pairs

Msg length 1024 bit (128B) Msg cp delay 336 (time unit) Msg chk delay 1 (time unit)		Effective Throughput with no Resend Pairs (Multicasting XTP V3.5) (TU = Time Unit) (1 sec = 10,000,000TU)			
Msg Inter (Time Unit)	Msg Rate (B/s)	No Delay		Delay	
		(B/s)	(KB/s)	(B/s)	(KB/s)
32,258	39,680	52,297	52.3	52,297	52.3
16,077	79,616	104,931	104.9	104,931	104.9
10,741	119,168	157,056	157.1	157,055	157.1
8,064	158,720	209,186	209.2	209,181	209.2
6,455	198,272	261,308	261.3	261,305	261.3
5,382	237,824	313,382	313.4	313,368	313.4
4,615	277,776	365,425	365.4	365,418	365.4
4,038	316,928	417,557	417.6	417,549	417.6
3,591	356,352	454,660	454.7	453,239	453.2
3,232	396,032	454,791	454.8	453,282	453.3
2,937	435,712	454,842	454.8	453,124	453.1
2,692	475,392	454,883	454.9	452,943	452.9
2,485	515,072	454,911	454.9	452,672	452.7

Table 11: Throughput vs Offered load for 128 byte messages with no Resend Pairs in Multicasting XTP V3.5

Msg length 1024 bit (128B) Msg cp delay 336 (time unit) Msg chk delay 1 (time unit)		Message Delay with no Resend Pairs (Multicasting XTP V3.5) (TU = Time Unit) (1 sec = 10,000,000TU)					
Msg Inter (TU)	Msg Rate (B/s)	No Delay			Delay		
		Min Delay (TU)	Msg Delay		Min Delay (TU)	Msg Delay	
			(TU)	(ms)		(TU)	(ms)
32,258	39,680	3591	2724.7	0.27	3592	2725.4	0.27
16,077	79,616	3591	2724.7	0.27	3591	2724.7	0.27
10,741	119,168	3591	2724.7	0.27	3591	2724.7	0.27
8,064	158,720	3591	2724.7	0.27	3591	2724.7	0.27
6,455	198,272	3591	2724.7	0.27	3591	2724.8	0.27
5,382	237,824	3591	2725.2	0.27	3591	2725.3	0.27
4,615	277,776	3591	2729.7	0.27	3590	2729.7	0.27
4,038	316,928	3592	2726.3	0.27	3591	2725.6	0.27
3,591	356,352	3595	2817.7	0.28	3591	2823.4	0.28
3,232	396,032	3596	3131.4	0.31	3591	3137.5	0.31
2,937	435,712	3596	3444.8	0.34	3590	3453.0	0.35
2,692	475,392	3596	3758.1	0.38	3591	3769.0	0.38
2,485	515,072	3595	4070.4	0.41	3591	4086.0	0.41

Table 12: Delay vs Offered load for 128 byte messages with no Resend Pairs in Multicasting XTP V3.5

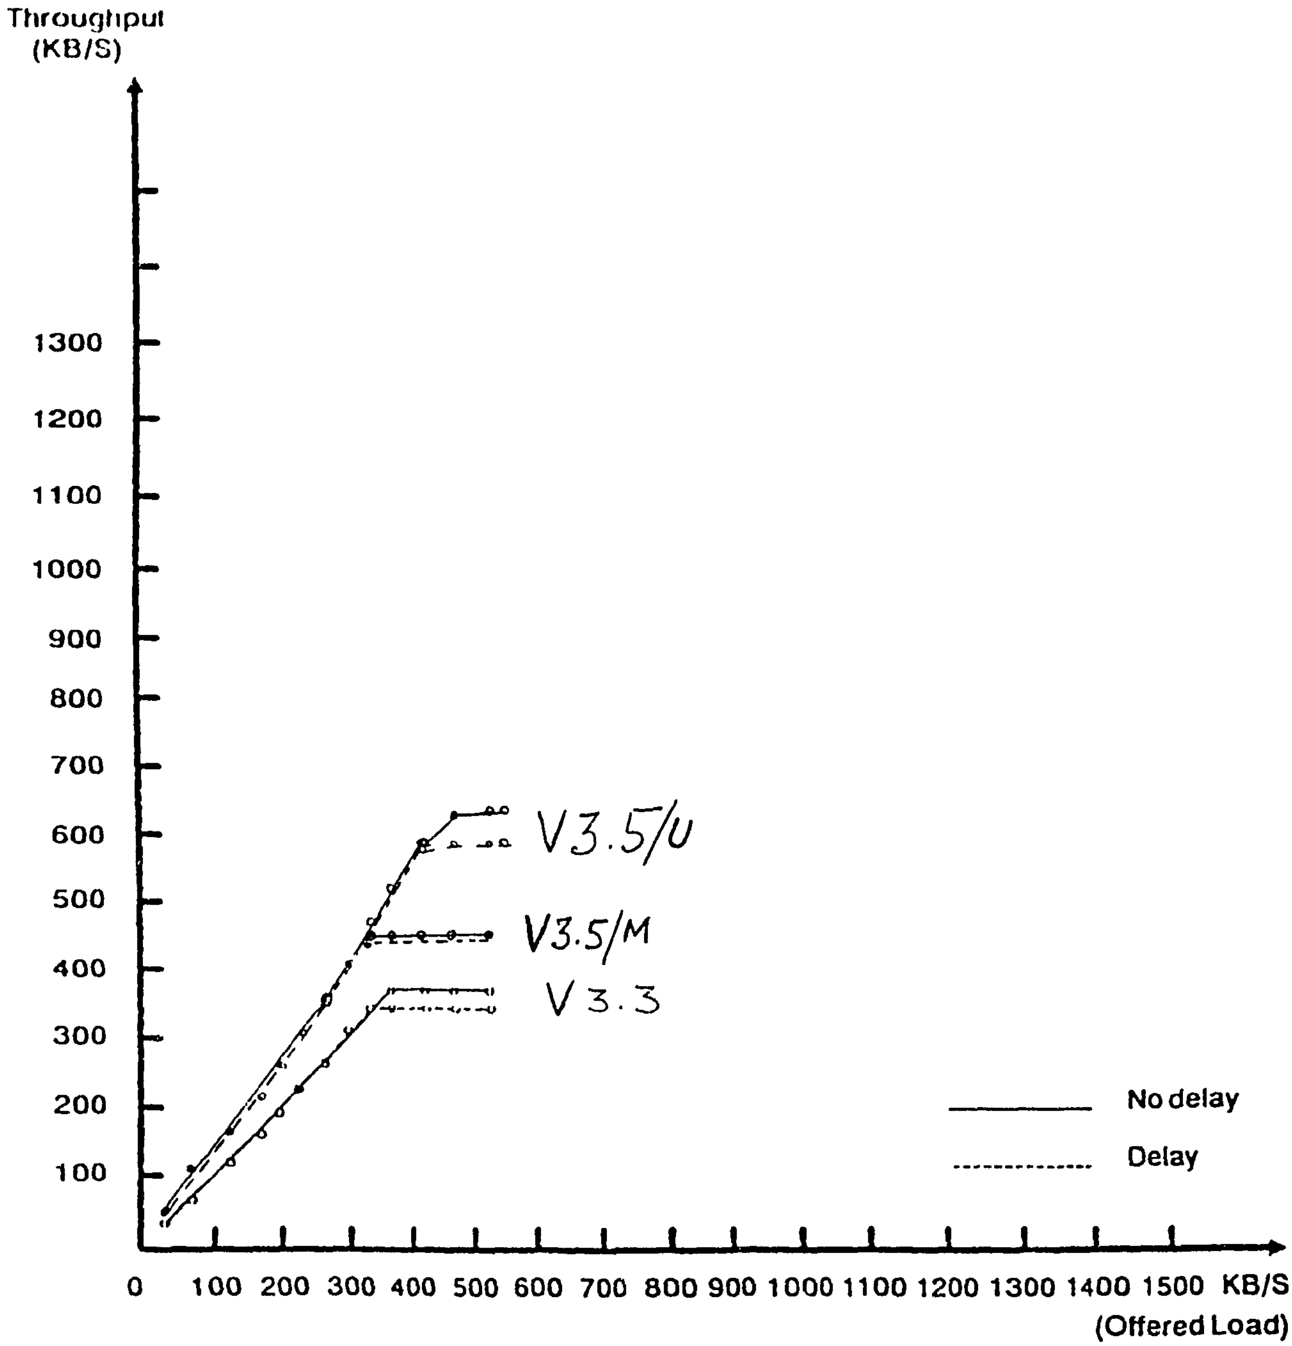

Figure 11: Throughput vs offered load for 128 byte messages with no Resend Pairs in Multicasting XTP V3.5

V3.5/M - Multicasting with no resend pairs
V3.5/U - Unicasting with no resend pairs
V3.3 - Unicasting with resend pairs

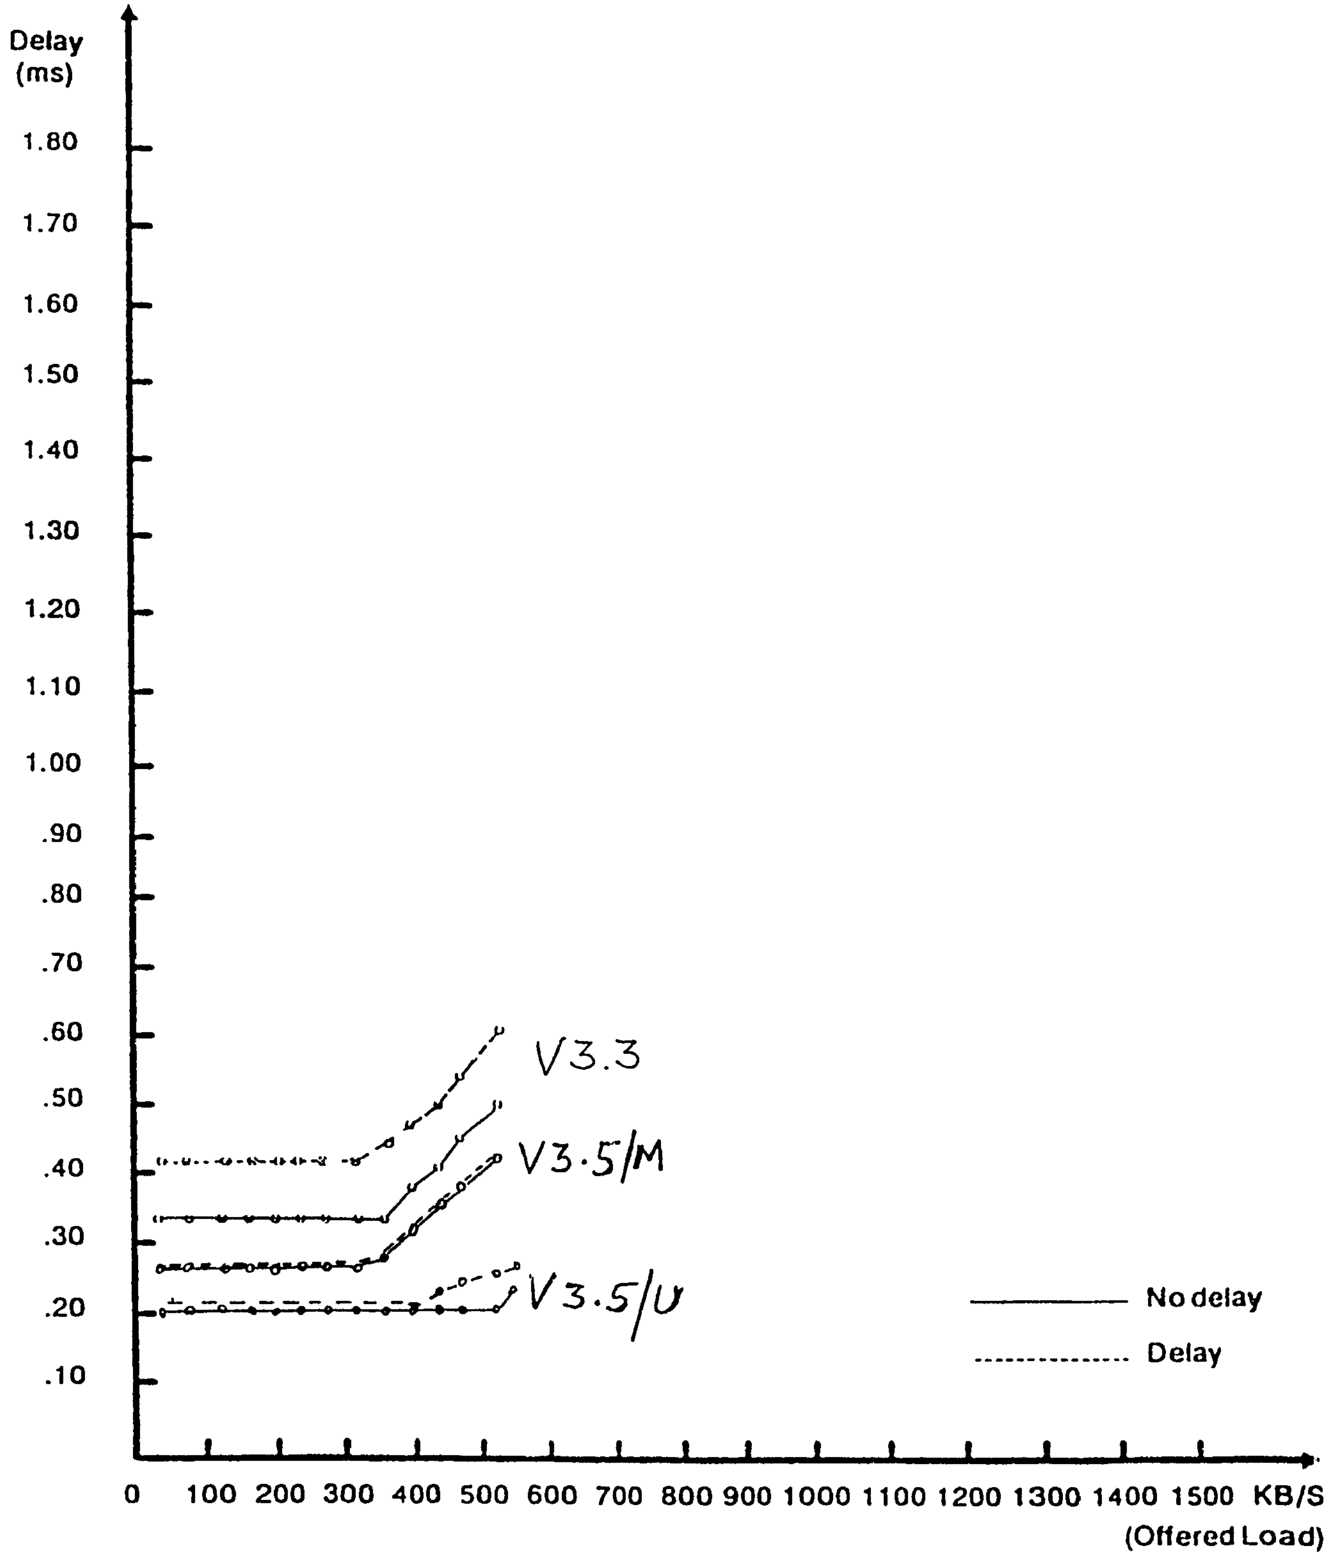

Figure 12: Delay vs offered load for 128 byte messages with no Resend Pairs in Multicasting XTP V3.5

V3.5/M - Multicasting with no resend pairs
V3.5/U - Unicasting with no resend pairs
V3.3 - Unicasting with resend pairs

Msg length 8192 bit (1KB) Msg cp delay 2,694 (time unit) Msg chk delay 1 (time unit)		Effective Throughput with no Resend Pairs (Multicasting XTP V3.5) (TU = Time Unit) (1 sec = 10,000,000TU)			
Msg Inter	Msg Rate	No Delay		Delay	
(Time Unit)	(B/s)	(B/s)	(KB/s)	(B/s)	(KB/s)
104,166	98,324	102,571	102.6	102,570	102.6
51,813	197,632	206,207	206.2	206,203	206.2
34,602	295,936	308,768	308.8	308,754	308.8
25,974	394,240	411,318	411.3	411,292	411.3
20,790	492,544	513,852	513.9	513,802	513.8
17,331	590,848	616,371	616.4	616,278	616.3
14,858	689,152	718,905	718.9	718,735	718.7
13,003	787,456	821,369	821.4	800,087	800.1
11,645	875,520	916,972	914.0	793,085	793.1
10,405	984,064	978,915	978.9	784,357	784.4
9,460	1,082,368	979,127	979.1	776,026	776.0
8,673	1,184,748	979,230	979.2	767,785	767.8
8,006	1,278,976	979,308	979.3	759,570	759.6

Table 13: Throughput vs Offered load for
1 Kbyte messages with no Resend Pairs
in Multicasting XTP V3.5

Msg length 8192 bit (1KB) Msg cp delay 21,559 (TU) Msg chk delay 1 (time unit)		Message Delay with no Resend Pairs (Multicasting XTP V3.5) (TU = Time Unit) (1 sec = 10,000,000TU)					
Msg Inter (TU)	Msg Rate (B/s)	No Delay			Delay		
		Min Delay (TU)	Msg Delay		Min Delay (TU)	Msg Delay	
			(TU)	(ms)		(TU)	(ms)
104,166	98,324	10,758	10313	1.03	12350	11839	1.18
51,813	197,632	10,758	10311	1.03	12350	11837	1.18
34,604	295,930	10,758	10311	1.03	12350	11837	1.18
25,974	394,240	10,759	10311	1.03	12350	11838	1.18
20,790	492,544	10,758	10312	1.03	12350	11839	1.18
17,331	590,848	10,759	10313	1.03	12350	11840	1.18
14,858	689,152	10,758	10313	1.03	12350	11842	1.18
13,003	787,456	10,759	10315	1.03	12350	12155	1.22
11,645	875,520	10,760	10274	1.03	12350	13634	1.36
10,405	984,064	10,762	10819	1.08	12350	15495	1.55
9,460	1082,368	10,762	11897	1.19	12350	17225	1.72
8,673	1184,748	10,765	13024	1.30	12350	19057	1.91
8,006	1278,976	10,765	14059	1.41	12350	20795	2.08

Table 14: Delay vs Offered load for 1 Kbyte messages with no Resend Pairs in Multicasting XTP V3.5

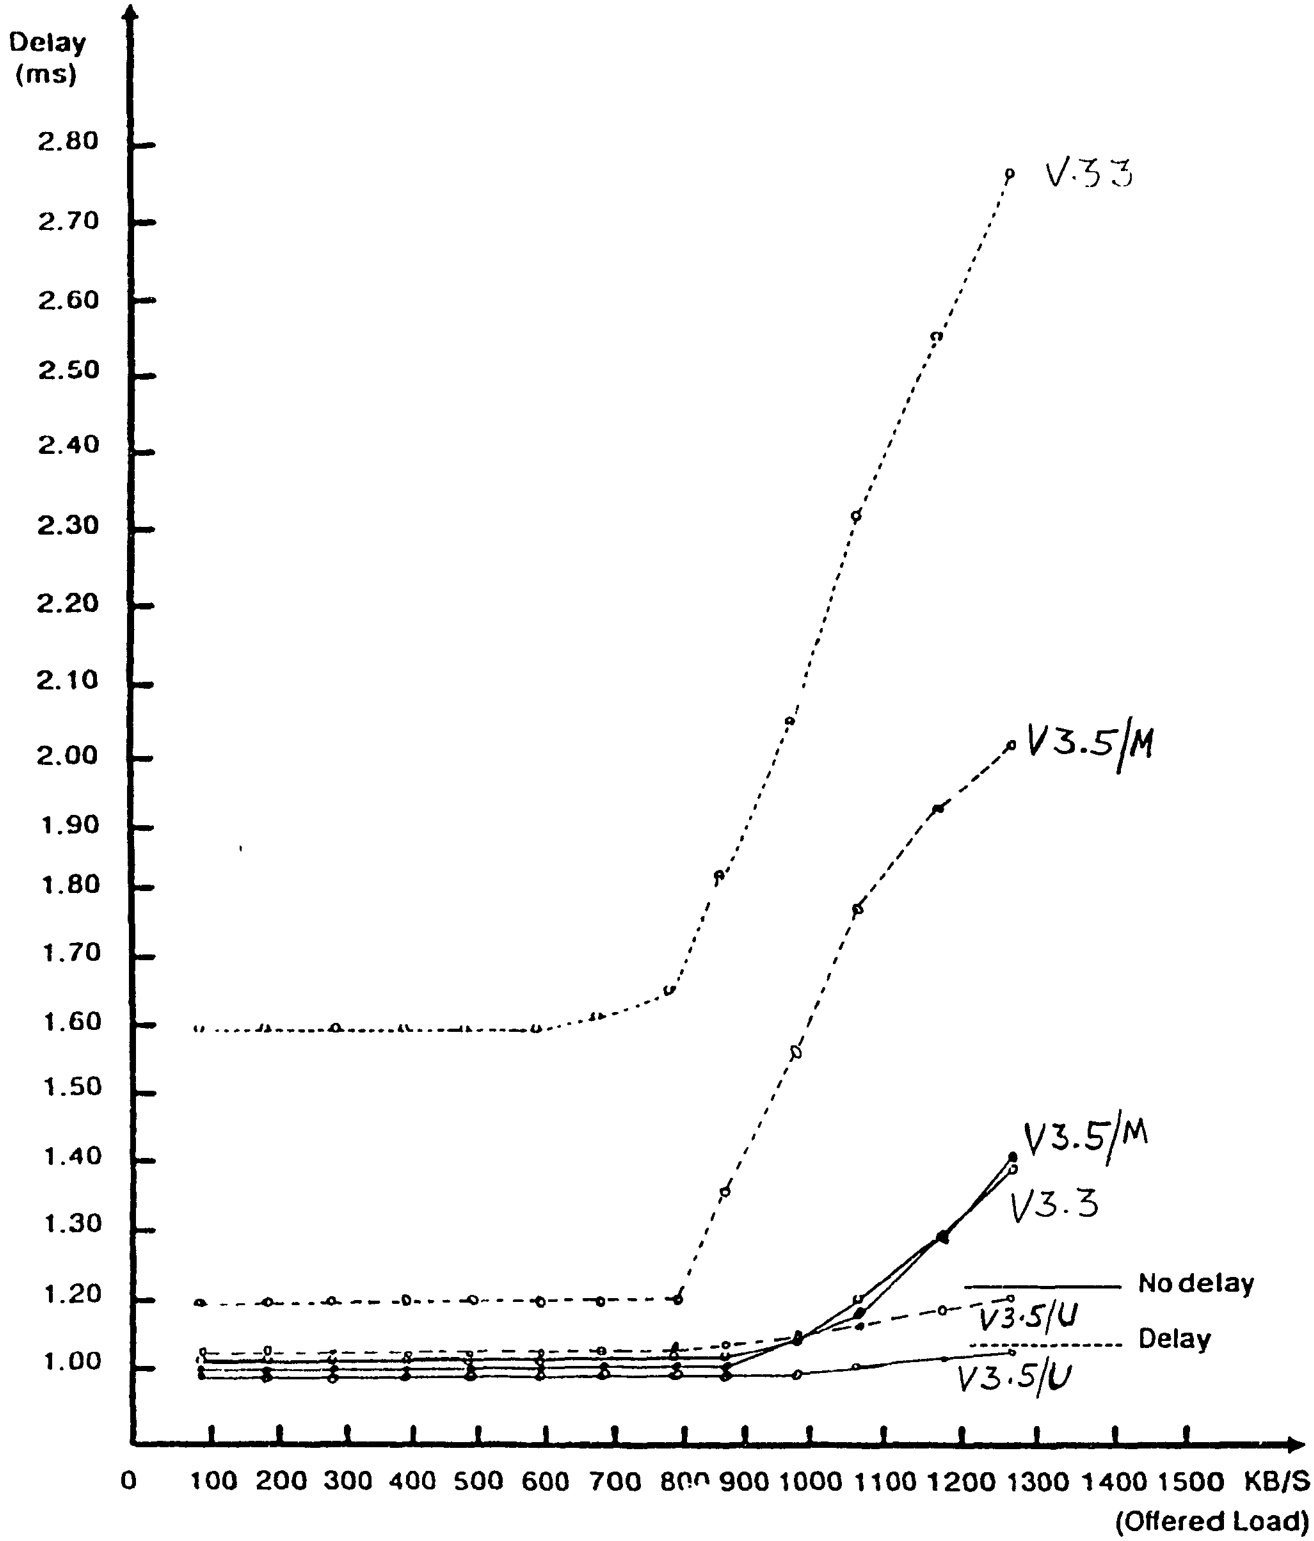

Figure 13: Throughput vs offered load for
1 Kbyte messages with no Resend Pairs
in Multicasting XTP V3.5

V3.5/M - Multicasting with no resend pairs
V3.5/U - Unicasting with no resend pairs
V3.3 - Unicasting with resend pairs

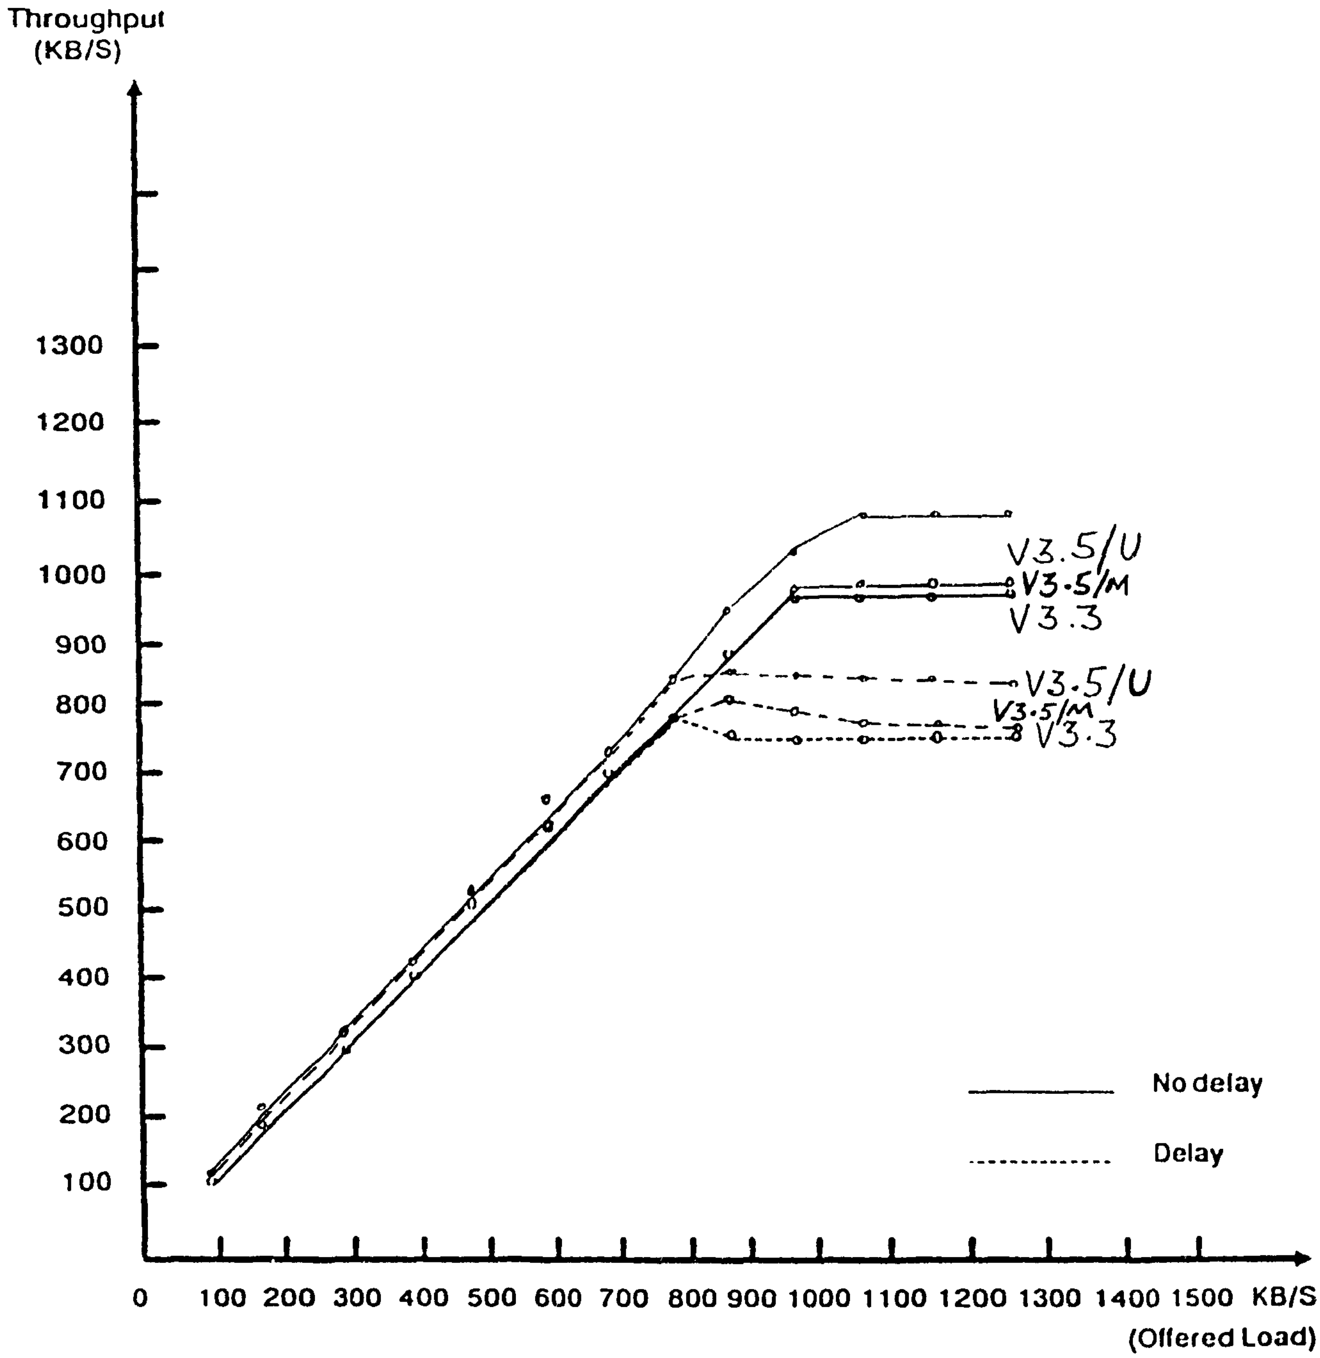

Figure 14: Delay vs offered load for
1 Kbyte messages with no Resend Pairs
in Multicasting XTP V3.5

V3.5/M - Multicasting with no resend pairs
V3.5/U - Unicasting with no resend pairs
V3.3 - Unicasting with resend pairs

Msg length 65,536 bit (8KB) Msg cp delay 21,559 (time unit) Msg chk delay 1 (time unit)		Effective Throughput with no Resend Pairs (Multicasting XTP V3.5) (TU = Time Unit) (1 sec = 10,000,000TU)			
Msg Inter	Msg Rate	No Delay		Delay	
(Time Unit)	(B/s)	(B/s)	(K/s)	(B/s)	(KB/s)
714,285	114,688	116,625	114.3	114,364	114.4
357,142	229,376	233,569	233.7	234,700	234.7
238,095	344,064	351,040	352.1	352,386	352.4
178,571	458,752	473,199	473.7	469,461	469.5
142,857	573,440	589,199	589.2	592,027	592.0
119,047	688,128	706,372	706.4	709,820	709.8
104,166	786,432	810,802	810.8	812,022	822.2
89,285	917,502	946,991	947.0	946,048	946.1
79,363	1,032,192	1,066,128	1066.1	1021,237	1021.2
71,428	1,146,880	1,144,554	1147.6	1013,896	1013.9
64,935	1,261,568	1,147,741	1147.7	1005,686	1005.7
59,524	1,376,256	1,147,893	1147.9	999,283	999.3
54,945	1,490,944	1,148,021	1148.0	990,920	990.9

Table 15: Throughput vs Offered load for 8 Kbyte messages With no Resend Pairs in Multicasting XTP V3.5

Msg length 65,536 bit (8KB) Msg cp delay 21,559 (TU) Msg chk delay 1 (time unit)		Message Delay with no Resend Pairs (Multicasting XTP V3.5) (TU = Time Unit) (1 sec = 10,000,000TU)					
Msg Inter (TU)	Msg Rate (B/s)	No Delay			Delay		
		Min Delay (TU)	Msg Delay		Min Delay (TU)	Msg Delay	
			(TU)	(ms)		(TU)	(ms)
714,285	114,688	9797	9634	0.96	10,948	10,979	1.10
357,142	229,376	9797	9621	0.96	10,948	10,700	1.07
238,095	344,064	9797	9602	0.96	10,948	10,690	1.07
178,571	458,752	9797	4498	0.95	10,948	10,698	1.07
142,857	573,440	9797	9535	0.95	10,948	10,604	1.06
119,047	688,128	9797	9544	0.95	10,948	10,613	1.06
104,166	786,432	9797	9503	0.95	10,948	9,101	0.91
89,285	917,502	9797	9492	0.95	10,948	9,836	0.98
79,363	1032,192	9796	9484	0.95	10,948	11,945	1.20
71,428	1146,880	9797	9791	0.95	10,948	12,384	1.24
64,935	1261,568	9797	10769	1.08	10,948	13,734	1.37
59,524	1376,256	9797	11746	1.18	10,948	15,078	1.51
54,945	1490,944	9796	12722	1.27	10,948	16,472	1.65

Table 16: Delay vs Offered load for
8 Kbyte messages with no Resend Pairs
in Multicasting XTP V3.5

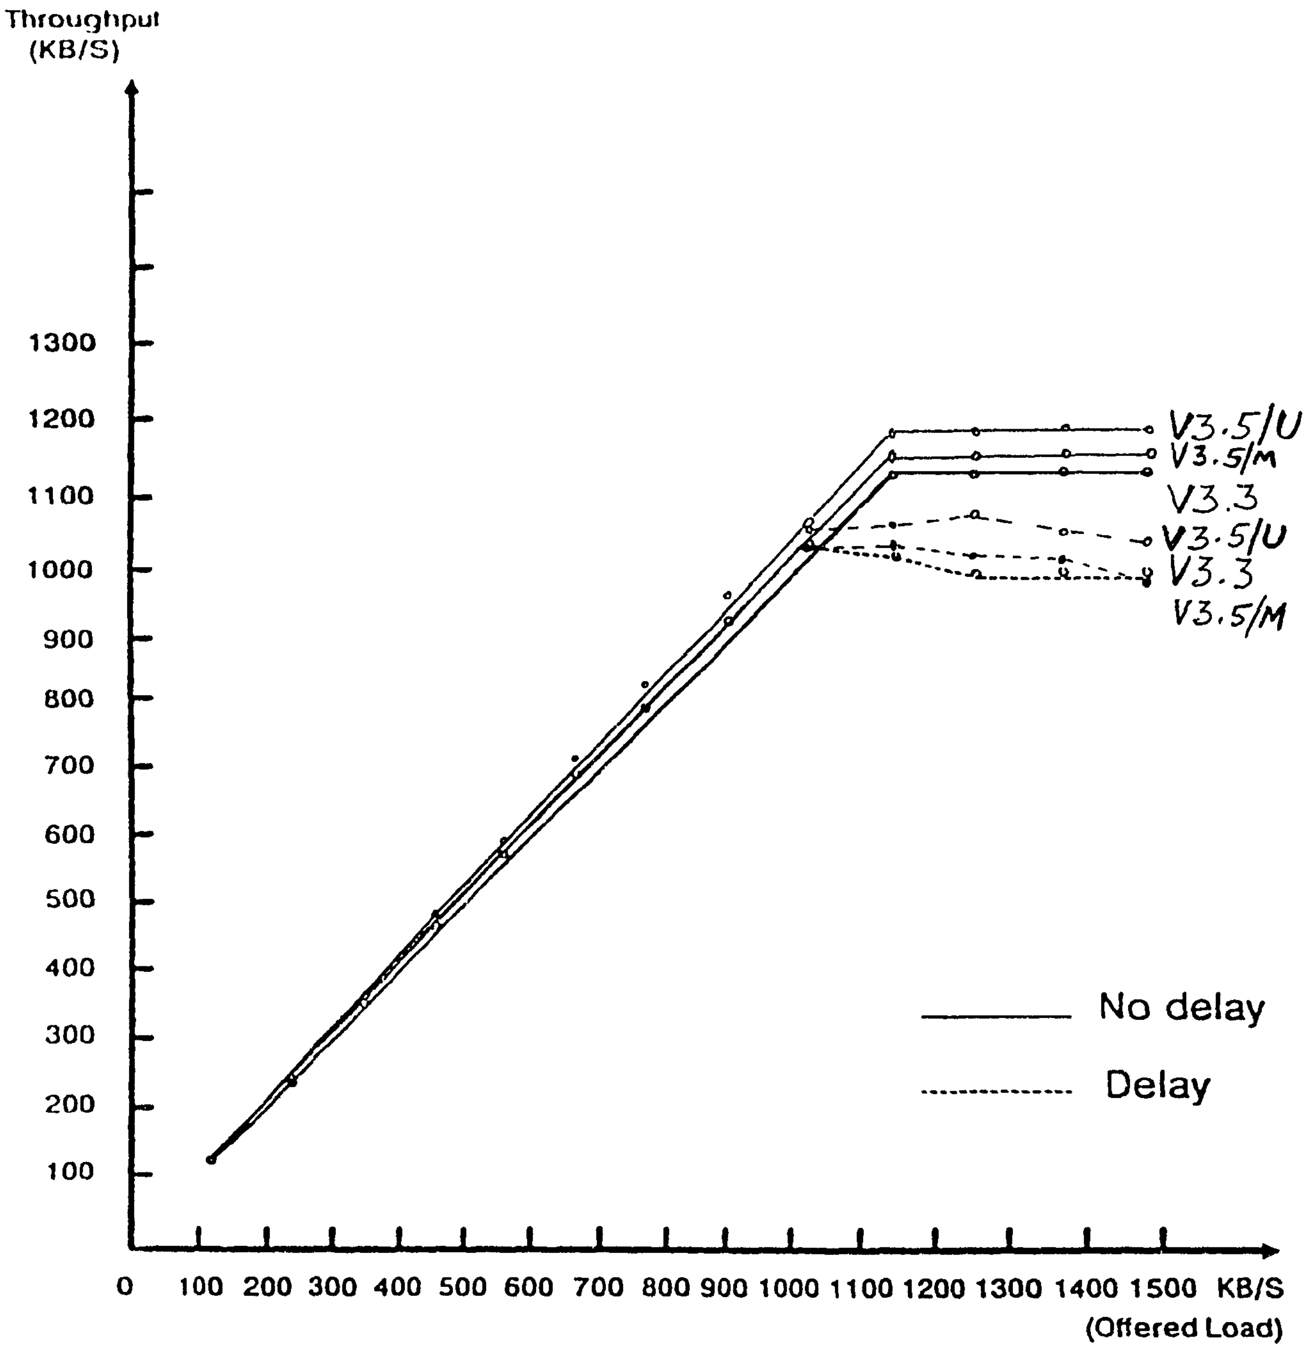

Figure 15: Throughput vs offered load for
8 Kbyte messages with no Resend Pairs
in Multicasting XTP V3.5

V3.5/M - Multicasting with no resend pairs
V3.5/U - Unicasting with no resend pairs
V3.3 - Unicasting with resend pairs

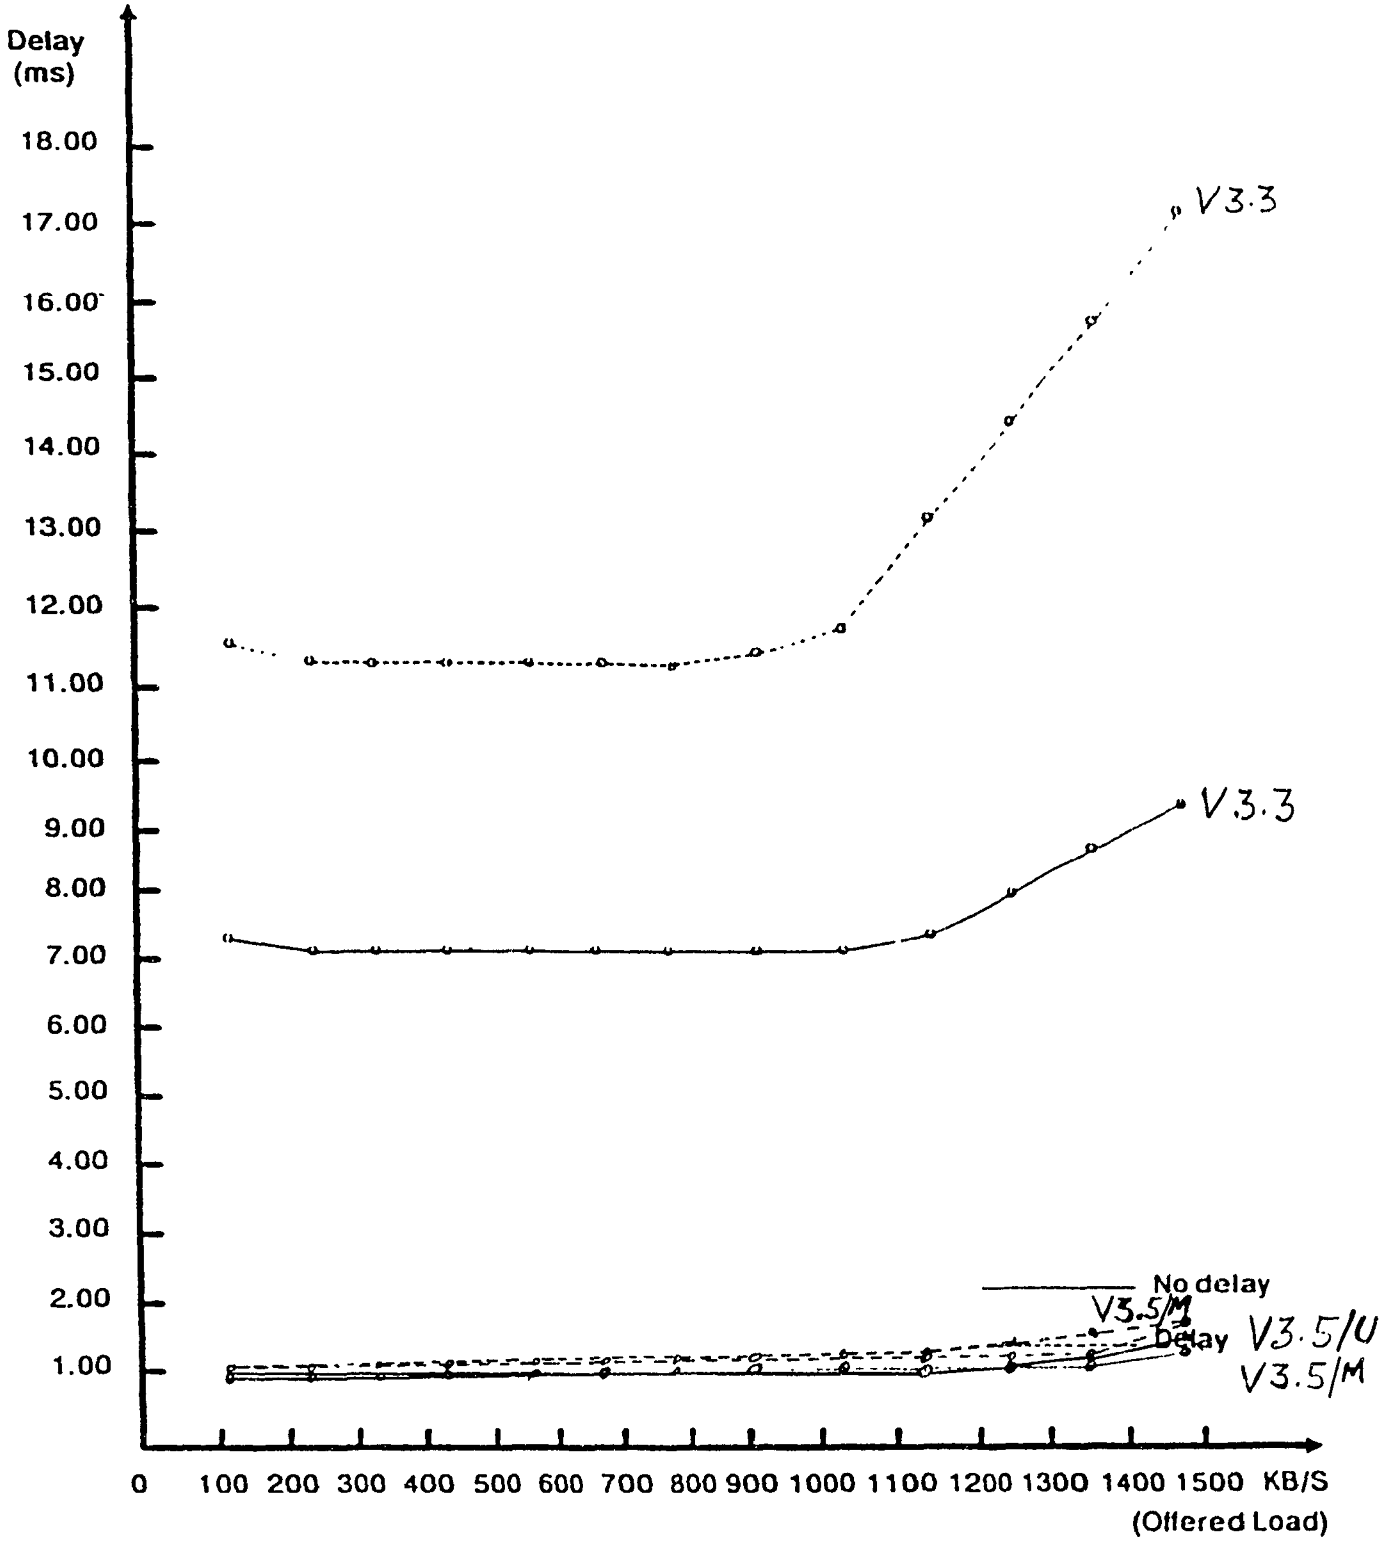

Figure 16: Delay vs offered load for 8 Kbyte messages with no Resend Pairs in Multicasting XTP V3.5

V3.5/M - Multicasting with no resend pairs
V3.5/U - Unicasting with no resend pairs
V3.3 - Unicasting with resend pairs

Message Size [bytes]	UNICASTING XTP V3.3 & V3.4 Maximum Expected Throughput [bytes/s]		
	Fixed Number of Resend Pairs [128 bytes]	No Resend Pairs [0 bytes]	Percent Improvement [%]
6	26,596	48,701	83.0
128	396,040	579,710	46.0
1024	984,615	1,092,150	10.9
8192	1,140,820	1,174,312	2.9

Table 17: Expected Throughput comparison with fixed and with no resend pairs in Unicasting XTP V3.3/V3.4

Message Size [bytes]	UNICASTING XTP V3.5 Maximum Expected Throughput [bytes/s]		
	Fixed Number of Resend Pairs [128 bytes]	No Resend Pairs [0 bytes]	Percent Improvement [%]
6	23,885	40,323	68.8
128	366,972	519,481	41.6
1024	990,712	1,063,123	7.3
8192	1,128,748	1,161,525	2.9

Table 18: Expected Throughput comparison with fixed and with no resend pairs in Unicasting XTP V3.5

Message Size [bytes]	MULTICASTING XTP V3.5 Maximum Expected Throughput [bytes/s]		
	Fixed Number of Resend Pairs [128 bytes]	No Resend Pairs [0 bytes]	Percent Improvement [%]
6	13,489	25,000	85.3
128	235,988	379,147	60.7
1024	813,215	971,168	19.4
8192	1,071,598	1,132,242	5.7

Table 19: Expected Throughput comparison with fixed and with no resend pairs in Multicasting XTP V3.5

Message Size [bytes]	**Maximum Expected Throughput Comparison** [bytes/s]		
	Unicasting V3.5 With Resend Pairs	**Multicasting** V3.5 With Resend Pairs	Multicasting = percent of Unicasting [%]
6	23,885	13,489	56.48
128	366,972	235,988	64.31
1024	990,712	813,215	82.08
8192	1,128,748	1,071,598	94.94

Table 20: Maximum Expected Throughput comparison with resend pairs in Unicasting and Multicasting XTP V3.5

Message Size [bytes]	Maximum Expected Throughput Comparison [bytes/s]		
	Unicasting V3.5 No Resend Pairs	**Multicasting** V3.5 No Resend Pairs	Multicasting = percent of Unicasting [%]
6	40,323	25,000	62.00
128	519,481	379,147	72.99
1024	1,063,123	971,168	91.35
8192	1,161,525	1,132,242	97.48

Table 21: Maximum Expected Throughput comparison with no resend pairs in Unicasting and Multicasting XTP V3.5

1Message Size [bytes]	Maximum Actual Throughput Comparison [bytes/s]		
	Unicasting V3.5 No Resend Pairs	**Multicasting** V3.5 No Resend Pairs	Multicasting = percent of Unicasting [%]
6	266,232	169,336	63.61
128	625,853	454,911	72.69
1024	1,080,019	979,308	90.68
8192	1,180,449	1,148,021	97.25

Table 22: Maximum Actual Throughput comparison with no resend pairs in Unicasting and Multicasting XTP V3.5

CHAPTER 5

CONCLUSION

5.1 General

Though multicasting simulation has been completed, it is difficult to prove the correctness of this implementation. It is only fair to say that if the XTP unicast implementation is correct, then the XTP Multicast implementation should also be correct.

5.2 XTP Multicasting Performance

The simulation indicates that XTP is capable of providing a highly efficient transport service to its users--up to 80% of the raw bandwidth usage in the file transfer applications in an ethernet environment. The XTP throughput performance also depends on how quickly data copying and checksumming can be done. The achievable throughput drops markedly when this factor is taken into account. In a no error environment, XTP performance for short packets (i.e., with no resend pairs) is improved by 83% for V3.3/V3.4, 69% for V3.5 (6 bytes packets) or 46% for V3.3/3.4, 42% for V3.5 (128 bytes packets) which is shown in tables 17 and 18. The resend pairs would have taken up 128 byte space which is a considerable amount of the extra overhead in the no error case.

The XTP V3.5 Multicasting performance with no resend pairs is about 63% of the unicasting performance (i.e., 50% of the raw bandwidth). One can also observe from tables 19, 20 and 22 that for larger message sizes, the performance increases. It appears that for large message sizes, the multicasting performance can achieve close to the unicasting performance.

Though the achievable throughput in multicasting is lower than unicasting, the advantage of time saving is in multicasting, where one sender can broadcast to many receivers (one-to-many), and onlv one of the receivers can respond when requested.

5.3 Future Work

As further research, the damping and slotting (when multicasting to a large number of receivers) and cloning (when large number of concentration channels are needed) techniques can be applied to XTP Multicast.

It would also be interesting to add buffer allocation measurement to the LANSF simulation. The count field in the queue structure and the maximum value field in the context could be used for this purpose.

REFERENCES

[1] Jason Xiao-Guang Chen: "Design and Validation of an XTP simulator", A Thesis in the Department of Computer Science, Concordia University, 1989.

[2] Gburzynki P. and Rudnicki P. "LANSF Protocol Modelling Environment, version 2.11e", Department of Computer Science, University of Alberta.

[3] Protocol Engines Incorporated: "XTP Protocol Definition, Revision 3.5", 10 September 1990.

[4] Stallings W.: "Data and Computer Communications", Macmillan Publishirg Co. Inc., U.S., 1985.

[5] Tannenbaum A. S.: "Computer Networks", Prentice-Hall Inc., New York, 1981.