Printing Requisition/Graphic Services

Dept. No.

97132

Title or Description CS-78-17

Date March 8/78 **Date Required** ASAP **Account** 126-6356-41

Signature Edith Huang **Signing Authority** Johnny W. Wong

Department Comp Science **Room** MC 5100B **Phone** 3402

Delivery ☐ Mail ☐ Pick-up ☐ Via Stores ☐ Other

1. Please complete unshaded areas on form as applicable. (4 part no carbon required).
2. Distribute copies as follows: White, Canary and Pink—Printing, Arts Library or applicable Copy Centre Goldenrod—**Retain.**
3. On completion of order, pink copy will be returned with printed material. Canary copy will be costed and returned to requisitioner, **Retain as a record of your charges.**
4. Please direct enquiries, quoting requisition number, to Printing/Graphic Services, Extension 3451.

Reproduction Requirements ☑ Offset ☐ Signs/Repro's ☐ Xerox

Number of Pages 37 **Number of Copies** 50

Type of Paper Stock ☑ Bond ☐ Book ☐ Cover ☐ Bristol ☐ Supplied

Paper Size ☑ 8½ x 11 ☐ 8½ x 14 ☐ 11 x 17

Paper Colour ☑ White ☐ Other

Ink ☐ Black

Printing ☑ 1 Side ☐ 2 Sides

Numbering to

Binding/Finishing Operations ☑ Collating ☐ Corner Stitching ☐ 3 Ring ☐ Tape ☐ Plastic Ring ☐ Perforating

Folding Finished Size

Cutting Finished Size

Special Instructions

Covers & backs supplied

Staple down left side

Cost: Time/Materials	Fun.	Prod.Un.	Prod. Cl.	Opr. No.	Mins.	Total
Signs/Repro's	1					
Camera	2					
Correcting & Masking Negatives	3					
Platemaking	4			25		
Printing	5			15		
Bindery	6			21-30 22.15		
Sub. Total Time						
Sub. Total Materials						
Prov. Tax						
Total						

Film		Plates	
Qty	Size	Qty	Size & Type

Paper		Plastic Rings	
Qty	Size	Qty	Size

Outside Services

Graphic Services May 68 482-2

TITLE: CS-78-17
Flow Control in Message-Switched Communication Networks

AUTHOR: J. Wong

NAME:	DATE TAKEN	DATE RETURNED	SIGNATURE
Dr. Karmeshu (Ext. 2736) Physics Dept.	Nov. 12, 1979.		

FLOW CONTROL IN MESSAGE-SWITCHED
COMMUNICATION NETWORKS

by

J. W. Wong

Research Report CS-78-17

Department of Computer Science and
Computer Communications Networks Group
University of Waterloo
Waterloo, Ontario, Canada

March 1978

Faculty

of

Mathematics

University of Waterloo
Waterloo, Ontario, Canada

FLOW CONTROL IN MESSAGE-SWITCHED COMMUNICATION NETWORKS

by

J. W. Wong

Research Report CS-78-17

Department of Computer Science and
Computer Communications Networks Group
University of Waterloo
Waterloo, Ontario, Canada

March 1978

This work was supported by the National Research Council and the Department of Communications of Canada.

To appear in Computer Communications.

SYNOPSIS

A queueing network model is used to analyse the performance of flow control techniques. Analytic expressions for throughput and mean end-to-end delay are derived. Numerical results show that under isarithmic control, the total network throughput is degraded significantly when the demand from one user is increased. Throughput degradation can be avoided if an end-to-end or a two-level control is used. Between these two schemes, the two-level control gives a higher throughput to the user with increased demand. This often results in a higher total throughput. The mean end-to-end delay under the various flow control techniques is also characterised.

INTRODUCTION

A message-switched communication network [1] is a collection of switching nodes connected together by a set of communication channels. It provides a message service to the users at the various nodes. Messages in this network are routed from one node to another in a store and forward manner until they reach their destination. As a message is routed through the network, it places demands on resources such as communication channels and buffer space in the switching nodes. We can thus view a message-switched communication network as a set of resources shared by a population of users. It is generally true that in resource-shared systems, if the resources are not properly managed, an increased demand due to a single user or a group of users may degrade the quality of service to the other users. In a message-switched communication network, this degradation can be in the form of increased end-to-end (i.e., source-to-destination) delay and/or reduced throughput [2-5]. There is therefore a need to control the flow of messages in the network. Of particular interest to this paper are admission control techniques where control is applied at the point of entry to the network.

A common admission control technique is the end-to-end control where a limit is placed on the number of messages in each virtual circuit. In a computer communication network, a virtual circuit [7] is a logical channel connecting together two users in the network. Examples of end-to-end control are the RFNM (Request For Next Message) feature in the ARPA network [4], and the window technique in DATAPAC [7].

A second technique for admission control is the isarithmic technique originally suggested by Davies [3]. This technique places a limit on the total number of messages in the network, no discrimination is made on the basis of source or destination. This is done by circulating a fixed number of "permits" in the network, and requiring a message to secure a permit before it can be admitted to the network. An isarithmic control scheme has been implemented by Price [6] in his network simulation model.

Davies [3] remarked that isarithmic control alone may not be effective, and suggested that it should be used as a supplement to other flow control techniques. This suggestion has motivated the development of a two-level isarithmic technique by Wong and Unsoy [9]. At the first level, a limit is placed on the total number of messages in the network. At the second level, disjoint groups of virtual circuits are formed and separate limits are placed on the number of messages belonging to each group. The limits at the second level not only provide a more effective flow control, they can also be adjusted to give a desirable number of permits to each group of virtual circuits.

The allocation of buffers to messages entering a switching node provides another method of admission control. As an example, Lam and Reiser [13] proposed an "input buffer limit" technique where messages entering a node are distinguished according to whether they are input messages or transit messages, and a limit is placed on the fraction of buffers that the input messages can occupy. The advantage of this technique is to give more buffers to messages already in transit so that network throughput can be improved. Irland [14] has also considered a model of a switching node with finite buffers. In his scheme, messages are distinguished by their out-going channels, and a limit is placed on the number of buffers allocated to each channel.

In this paper, we will use a queueing network model of the type analysed by Baskett, et.al. [10] to study the performance characteristics of admission control techniques. As pointed out by Lam and Reiser [13] and Irland [14], the queueing analysis of a network model with finite buffers at each node is very difficult. We will thus assume that the buffer space is infinite and restrict our analysis to the isarithmic, end-to-end, and two-level control. Approximation techniques to analyse network models with finite buffer space can be found in [13,14].

Pennotti and Schwartz [2] have analysed the end-to-end flow control scheme for a store and forward tandem link. Their work was later extended by Chatterjee, et.al. [8] to include random routing. The model analysed in this paper is different from those found in [2,8] in the sense that it is a total network model, and all source-destination pairs are taken into consideration.

In what follows, we will describe the basic network model, derive analytic expressions for throughput and mean end-to-end delay, and present numerical examples to compare the performance of the isarithmic, end-to-end, and two-level control schemes.

MODEL DESCRIPTION

Our model is essentially a generalisation of Kleinrock's classical model [1] to include flow control. Emphasis is placed on the distinction of messages by their source and destination. We first assume, as in [1], that the delay experienced by a message in a message-switched network is approximated by the queueing time and data transfer time in the channels. The processing time at the switching nodes and the propagation delays are assumed to be negligible. Let M be the number of channels and C_i be the capacity of channel i, $i = 1,2,\ldots,M$. In our queueing network model, each of the M channels is represented by a single server queue. The queueing discipline at each channel is first-come, first-served. We assume that all channels are error-free and all nodes have unlimited buffer space.

Messages are classified according to source-destination pairs. In particular, a message is said to belong to class (s,d) if its source node is s and its destination node is d. Let R be the total number of message classes. In a network with N switching nodes, R can be as large as N(N-1). For convenience, we assume that message classes are numbered from 1 to R, and we use r instead of (s,d) to denote a message class. The arrival process of class r messages from outside the network is assumed to be Poisson with mean rate

γ - gamma

$\gamma(r)$. Message lengths for all classes are assumed to have the same exponential distribution, and we use $1/\mu$ to denote the mean message length. It follows from this last assumption that the data transfer time of all messages at

μ - mu

channel i is exponential with mean $1/(\mu C_i)$. For the mathematical analysis to be tractable, Kleinrock's independence assumption [1] is used. This assumption states that each time a message enters a switching node, a new length is chosen from the exponential message length distribution.

For convenience, we assume that a fixed routing algorithm is used. In fixed routing, a unique path (or ordered set of channels) is defined for each source-destination pair. Our results can easily be generalised to include random routing.

To model isarithmic control, we assume that there is no delay in circulating "permits" through the network, and permits are allocated to messages arriving from outside the network on a first-come, first-served basis, no discrimination is made on the basis of source or destination. A similar assumption is made for the two-level control, where a message must acquire two types of permits (one for each level) before it can enter the network.

To model end-to-end control, we assume that a single limit is placed on the number of messages in all virtual circuits belonging to the same source-destination pair. We thus have separate limits on the number of messages belonging to each class. This abstraction allows us to simplify our model while keeping the fundamental feature of end-to-end control.

To complete our model for admission control, we assume that the message classes are divided into D disjoint groups and a message is said to belong to group u (denoted by G_u) if its class number is in group u. For a state S of our network model, let $|S|_u$ be the number of group u messages and $|S|$ be the total number of messages (from all groups) in the network. We define the following limits for our two-level control scheme:

First level: $|S| \le L$

Second level: $|S|_u \le L_u$ for $u = 1,2,\ldots,D$

With these limits, a group u message arriving from outside the network is not allowed to enter the network if $|S| = L$ or $|S|_u = L_u$. We assume that such arrivals are turned away and will not return (i.e., lost).

It is easy to see that the isarithmic and end-to-end control schemes are both special cases of this two-level control. They are given by the following limits:

(1) Isarithmic: there is only one group with $L_1 = L$.

(2) End-to-End: there are R groups, each group has a single source-destination pair, and $L = L_1 + L_2 + \ldots + L_R$.

QUEUEING ANALYSIS

EQUILIBRIUM STATE PROBABILITIES

We first consider a state of our network model given by $S = (S_1, S_2, \ldots, S_M)$ where $S_i = (n_{i1}, n_{i2}, \ldots, n_{iR})$ and n_{ir} is the number of class r messages at channel i. A feasible state is characterised by $\sum_{i=1}^{M} \sum_{r=1}^{R} n_{ir} \le L$ and $\sum_{i=1}^{M} \sum_{r \in G_u} n_{ir} \le L_u$. We also require that $n_{ir} = 0$ if class r messages are not routed through channel i. To get the equilibrium state probabilities, we follow the solution technique of Baskett, et.al. [10] and find that:

$$P(S = (S_1, S_2, \ldots, S_M)) = K \prod_{i=1}^{M} \left[\sum_{r=1}^{R} n_{ir} \right]! \prod_{r=1}^{R} (1/n_{ir}!)[\lambda_{ir}/(\mu C_i)]^{n_{ir}} \qquad (1)$$

λ - lambda where λ_{ir} is the total mean arrival rate of class r messages to channel i conditioned on no message being lost. Since we have assumed a fixed routing algorithm, λ_{ir} is given by:

$$\lambda_{ir} = \begin{cases} \gamma(r) & \text{if class r messages are routed through channel i} \\ 0 & \text{otherwise} \end{cases}$$

K is a normalisation constant obtained by summing the probabilities of all feasible states and equating the sum to one.

A complete derivation of Eq.(1) can be found in [9].

We next define a less detailed state description given by $S = (y_1, y_2, \ldots, y_M)$ where $y_i = (m_{i1}, m_{i2}, \ldots, m_{iD})$ and m_{iu} is the number of group u messages at channel i. A feasible state is now characterised by $\sum_{i=1}^{M} \sum_{u=1}^{D} m_{iu} \le L$, $\sum_{i=1}^{M} m_{iu} \le L_u$, and $m_{iu} = 0$ if group u messages are not routed through channel i. By summing all state probabilities $P(S = (S_1, S_2, \ldots, S_M))$ such that $\sum_{r \varepsilon G_u} n_{ir} = m_{iu}$ for $i = 1,2,\ldots,M$, and $u = 1,2,\ldots,D$, we get [9]:

$$P(S = (y_1, y_2, \ldots, y_M)) = K \prod_{i=1}^{M} \left[\sum_{u=1}^{D} m_{iu} \right]! \prod_{u=1}^{D} (1/m_{iu}!) [\xi_{iu}/(\mu C_i)]^{m_{iu}} \qquad (2)$$

where

ξ - xi $$\xi_{iu} = \sum_{r \varepsilon G_u} \lambda_{ir}$$

The normalisation constant K can be computed from:

$$K = \left[\sum_{\substack{\text{all feasible} \\ (y_1, y_2, \ldots, y_M)}} \prod_{i=1}^{M} \left[\sum_{u=1}^{D} m_{iu} \right]! \prod_{u=1}^{D} (1/m_{iu}!) [\xi_{iu}/(\mu C_i)]^{m_{iu}} \right]^{-1} \quad (3)$$

PERFORMANCE MEASURES -- MESSAGE GROUPS

We now derive the expressions for throughput and mean end-to-end delay of each message group. let $Q(m_1, m_2, \ldots, m_D)$ be the equilibrium probability that the number of group u messages in the network is m_u, $u = 1, 2, \ldots, D$.

$$Q(m_1, m_2, \ldots, m_D) = \sum_{m_{1u}+m_{2u}+\ldots+m_{Mu} = m_u, \text{ all } u} P(S = (y_1, y_2, \ldots, y_M)) \quad (4)$$

The external arrival rate of group u messages is:

$$\gamma_u = \sum_{r \in G_u} \gamma(r)$$

Since a group u message is not lost if $m_u < L_u$ and $m_1+m_2+\ldots,+m_D < L$, we have the following expression for group u throughput:

$$\gamma_u^* = \gamma_u \sum_{\substack{m_u < L_u \text{ and} \\ m_1+m_2+\ldots+m_D < L}} Q(m_1, m_2, \ldots, m_D)$$

The mean number of group u messages in the network is given by:

$$\bar{m}_u = \sum_{k=0}^{L_u} k \sum_{\substack{\text{all feasible} \\ (m_1,m_2,\ldots,m_D) \\ \text{with } m_u = k}} Q(m_1,m_2,\ldots,m_D)$$

where a feasible $(m_1,m_2,\ldots,m_D)$ is characterised by $m_u \le L_u$, $u = 1,2,\ldots,D$ and $m_1+m_2+\ldots+m_D \le L$. Finally, we apply Little's result [11] and get the following expression for the mean end-to-end delay of group u messages:

$$\bar{T}_u = \bar{m}_u / \gamma_u^*$$

We see from the above derivation that before we can compute numerical values for the performance measures, we must first compute the normalisation constant K. An efficient algorithm to compute K and $Q(m_1,m_2,\ldots,m_D)$ is given in the Appendix.

PERFORMANCE MEASURES -- MESSAGE CLASSES

We now consider the throughput and mean end-to-end delay of each message class. Let $r \in G_u$. It is obvious that if class r is the only member of group u, the expressions derived in the last section are directly applicable. If class r is not the only member of group u, then we have the following expression for class r throughput:

$$\gamma^*(r) = \gamma_u^* \, \gamma(r) \, / \, \gamma_u$$

To derive the mean end-to-end delay of class r messages, we divide G_u into two artificial groups: $G_{u,1}$ containing class r alone and $G_{u,2}$ containing the remaining classes in G_u. We also add the following limits:

$$|S|_{u,1} \le L_u$$

$$|S|_{u,2} \le L_u$$

$$|S|_{u,1} + |S|_{u,2} \le L_u$$

This technique allows us to isolate class r while maintaining the limits for group u. It can be shown quite easily that the solution to the equilibrium state probabilities in Eq.(2) can be extended to include these new limits. We can thus use the derivations in the last section to get $\overline{T}_{u,1}$ which gives us the mean end-to-end delay of class r messages.

As a final remark, the same technique can be used to get the throughput and mean end-to-end delay of messages belonging to a sub-group.

NUMERICAL EXAMPLES AND DISCUSSION OF RESULTS

In this section, we present numerical examples to illustrate the performance characteristics of the isarithmic, end-to-end, and two-level control schemes. We first note from the Appendix that the amount of computation is on the order of $MDL_1 L_2 \ldots L_D$. It is easy to see that this amount grows exponentially with D, and for networks with a reasonably large number of message groups, the computation of performance measures may be practically impossible. Our example models are therefore designed in such a way that the amount of computation is not excessive, but are sufficiently complex to illustrate the performance characteristics of the three flow control techniques.

EXAMPLE 1 -- TANDEM LINK

Our first example is based on the tandem link model shown in Figure 1. It has 4 nodes and 3 channels. The channel capacities are selected such that the mean data transfer time at channels 1 and 2 is 0.1 and that at channel 3 is 0.2. We assume that all external arrival rates are zero except for two message classes: class 1 with (s,d) = (1,3) and class 2 with (s,d) = (2,4). These two classes will therefore compete for the use of channel 2. We assume that $\gamma(1) = 2$ and $\gamma(2) = 4$ correspond to the normal demand by class 1 and 2 respectively, and that the two message classes are members of separate groups. In particular, we have $G_1 = \{(1,3)\}$ and $G_2 = \{(2,4)\}$. The following three schemes are considered:

(A) Isarithmic with $L = 8$ and $L_1 = L_2 = 8$

(B) Two-level control with $L = 8$ and $L_1 = L_2 = 6$

(C) End-to-End with $L = 8$ and $L_1 = L_2 = 4$

We first fix $\gamma(1)$ at 2 (normal class 1 demand) and study the effect of an increase in $\gamma(2)$ (class 2 demand) on network performance. The throughput of each class is shown in Figure 2. We observe that isarithmic control gives the highest level of class 2 throughput. The level of class 1 throughput, however, is much reduced. As a consequence, the total throughput increases at first, and then declines considerably. This behaviour is shown in Figure 3. With end-to-end or two-level control, the reduction of class 1 throughput is much smaller, and the level of class 1 throughput is also insensitive to changes in $\gamma(2)$. Between these two schemes, the two-level control performs better in terms of total throughput because class 2 messages can have up to 6 permits (rather than 4 in end-to-end control) while 2 permits are sufficient to achieve a reasonable level of class 1 throughput.

We thus conclude that isarithmic control is not capable of preventing throughput degradation when the demand from one user is increased. Throughput degradation can be avoided if an end-to-end or a two-level control scheme is used. Between these two schemes, two-level control allows a user in heavy demand to have more permits. This often results in a higher total throughput. These conclusions are consistent with Davies' suggestion [3] that isarithmic control should be used as a supplement to other flow control schemes.

The mean end-to-end delay of each class is plotted in Figure 4. We observe that when $\gamma(2)$ is large, there is no increase in class 1 delay, and a control scheme with a larger L_2 results in a higher class 2 delay.

We next fix $\gamma(1)$ at 2.0 (normal class 1 demand) and $\gamma(2)$ at 12.0 (heavy class 2 demand) and study the performance of end-to-end control as a function of L_2 (with $L_1 = L_2$). The results are plotted in Figures 5 and 6. We observe that the throughput and mean end-to-end delay of class 1 messages are not affected significantly. As to class 2 messages, the throughput is a slowly increasing function of L_2 while the mean end-to-end delay increases almost linearly with L_2. These observations are not surprising because class 2 throughput is limited by the capacity of the slowest channel; and each time L_2 is increased by 1, the mean number of class 2 message waiting at the slowest channel is increased by approximately 1. The mean end-to-end delay is therefore increased by approximately the mean data transfer time at this channel. A good strategy is then to select L_2 such that the throughput is relatively high and class 2 delay is within some acceptable limit.

Finally, we consider a family of two-level control schemes with L = 8. $\gamma(1)$ and $\gamma(2)$ are still fixed at 2.0 and 12.0 respectively. L_1 and L_2 (with $L_1 = L_2$) are allowed to vary between 4 and 8. The results are shown in Figures 7 and 8. We observe that the total throughput reaches a maximum at $L_1 = L_2 = 6$ and then declines. One can therefore adjust the limits at the second level so that the throughput is highest. As to the mean end-to-end delay, we have the same properties as those observed in end-to-end control, i.e., class 1 delay is not affected significantly, and class 2 delay increases linearly with L_2.

EXAMPLE 2 -- TOTAL NETWORK MODEL

Our second example is based on the network model shown in Figure 9. This network has 5 nodes and 10 channels. All channels are assumed to have the same capacity and the mean message length is chosen such that the mean data transfer time at each channel (i.e., $1/(\mu C_i)$) has a value of 0.1. The routing algorithm is based on the shortest path. In our example network, there is a unique shortest path between each pair of nodes. We assume that there are two message groups. Group 1 contains the source-destination pairs among nodes 1, 2, and 3, i.e.,

$$G_1 = \{(1,2),(1,3),(2,1),(2,3),(3,1),(3,2)\}.$$

Group 2 includes all source-destination pairs not in group 1, i.e.,

$$G_2 = \{(s,d) \mid (s,d) \notin G_1\}.$$

The external arrival rate of messages belonging to each source-destination pair is given by the traffic matrix in Figure 10. These rates are expressed in terms of α and β which can vary to reflect the demand put on the network by each message group. The case $\alpha = \beta = 1.0$ corresponds to the normal demand.

For convenience, we use the vector (L, L_1, L_2) to denote the two-level control with two message groups. The following four schemes are considered:

(D) (15,15,15) -- isarithmic control

(E) (15,12,6) -- two-level control with $L_1/L_2 = 2.0$

(F) (15,9,9) -- two-level control with $L_1/L_2 = 1.0$

(G) (15,6,12) -- two-level control with $L_1/L_2 = 0.5$

We fix α at 1.0 (normal demand from group 1) and allows β (group 2 demand) to increase. The results are plotted in Figures 11 to 14. We observe once again that isarithmic control is not capable of preventing throughput degradation when the demand of one group is increased. If a two-level control is used, the total throughput is much improved, and the level of total throughput is also not sensitive to variations in β.

The plots in Figures 11 and 12 also show the effect of the ratio L_1/L_2 on the throughput of each group. In particular, the level of group 1 (or 2) throughput is higher if a larger (or smaller) ratio is used for L_1/L_2. The limits L, L_1, and L_2 can therefore be adjusted to give preferential treatment to a particular message group in terms of a higher throughput. Similar properties are also found in end-to-end control.

As to the mean end-to-end delay, we have essentially the same observation as that made from Figure 4, namely, group 1 delay is not increased, and group 2 delay increases with L_2.

ACKNOWLEDGEMENT

This work was supported in part the National Research Council and the Department of Communications of Canada.

REFERENCES

1 Kleinrock, L Communication Nets -- Stochastic Message Flow and Delays McGraw-Hill, New York (1964)

2 Pennotti, M and Schwartz, M 'Congestion Control in Store and Forward Tandem Links' IEEE Trans. on Comm. (December 1975) pp 1434-1443

3 Davies, D 'The Control of Congestion in Packet Switching Networks' IEEE Trans. on Comm. (June 1972) pp 546-550

4 Kahn, R and Crowther, W 'Flow Control in Resource Sharing Computer Network' IEEE Trans. on Comm. (June 1972) pp 539-546

5 Rudin, H 'Flow-Control: Session Chairman's Remarks' Proc. International Conference on Computer Communications, Toronto, Canada (August 1976) pp 463-466

6 Price, W 'Simulation of a Packet-Switched Data Network Operating under Isarithmic Control with Revised Link and Node Protocol' NPL Report COM 71 (September 1973)

7 Twyver, D and Rybczynski, A "DATAPAC Subscriber Interfaces" Proc. International Conference on Computer Communications, Toronto, Canada (August 1976) pp 143-149

8 Chatterjee, A Georganas, N and Verma, P 'Analysis of a Packet Switched Network with End-to-End Congestion Control and Random Routing' Proc. International Conference on Computer Communications, Toronto, Canada (August 1976) pp 488-494

9 Wong, J and Unsoy, M 'Analysis of Flow Control in Computer Networks' CCNG Report E-54, University of Waterloo, Canada (November 1976)

10 Baskett, F Chandy, K Muntz, R and Palacios, F 'Open, Closed and Mixed Networks of Queues with Different Classes of Customers' Journal of the ACM (April 1975) pp 248-260

11 Little, J 'A proof of the Queueing Formula $L = \lambda W$' Operations Research (1961) pp 383-387

12 Reiser, M and Kobayashi, H 'On the Convolution Algorithm for Separable Queueing Networks' Proc. International Symposium on Computer Performance Modelling, Harvard University (March 1976) pp 109-117

13 Lam, S and Reiser, M 'Congestion Control of Store-and-Forward Networks by Input Buffer Limits' to appear in Proc. National Telecommunications Conference, Los Angeles (December 1977)

14 Irland, M 'Analysis and Simulation of Congestion in Packet-Switched Networks' CCNG Report E-61, University of Waterloo, Canada (April 1977)

APPENDIX: COMPUTATION OF K AND $Q(m_1, m_2, \ldots, m_D)$

The normalisation constant K, as given by Eq.(3), can be computed from:

$$K^{-1} = \sum_{\substack{\text{all feasible} \\ (y_1, y_2, \ldots, y_M)}} F_1(y_1)F_2(y_2)\ldots F_M(y_M)$$

where $F_i(y_i) = \left[\sum_{u=1}^{D} m_{iu}\right]! \prod_{u=1}^{D} (1/m_{iu}!)[\xi_{iu}/(\mu C_i)]^{m_{iu}}$

The algorithm to compute K is based on the work of Reiser and Kobayashi [12]. Let $y_i^* = (m_{i1}^*, m_{i2}^*, \ldots, m_{iD}^*)$ where $m_{iu}^* = \sum_{j=1}^{i} m_{ju}$ is the total number of group u messages in channels 1, 2, ..., i-1 and i. A feasible y_i^* is then characterised by $\sum_{j=1}^{i} \sum_{u=1}^{D} m_{ju}^* \leq L$ and $\sum_{j=1}^{i} m_{ju}^* \leq L_u$ for $u = 1,2,\ldots,D$. We also define:

$$F_i^*(y_i^*) = \sum_{y_1+y_2+\ldots+y_i = y_i^*} F_1(y_1)F_2(y_2)\ldots F_i(y_i)$$

It is easy to see that $F_1^*(y_1^*) = F_1(y_1^*)$, and for $i > 1$, $F_i^*(y_i^*)$ can be computed from the following recursive formula [12]:

$$F_i^*(y_i^*) = F_{i-1}^*(y_i^*) + \sum_{u=1}^{D} F_i^*(y_i^* - e_u)[\xi_{iu}/(\mu C_i)]$$

where

$$e_u = (0, \ldots, 0, \overset{\substack{\text{u-th}\\ \downarrow}}{1}, 0, \ldots, 0)$$

K can then be computed from:

$$K^{-1} = \sum_{\text{all feasible } y_M^*} F_M^*(y_M^*)$$

The amount of computation is $\mathcal{O}(MDL_1L_2\ldots L_D)$, and the storage requirement is $\mathcal{O}(L_1L_2\ldots L_D)$.

To compute $Q(m_1,m_2,\ldots,m_D)$, we get from Eq.(4) that $Q(m_1,m_2,\ldots,m_D)$ can be rewritten as:

$$Q(m_1,m_2,\ldots,m_D) = K\, F_M^*(y_M^*) \Big|_{y_M^* = (m_1,m_2,\ldots,m_D)}$$

which can be computed directly from the $F_M^*(y_M^*)$'s.

Figure 1 Tandem Link Model

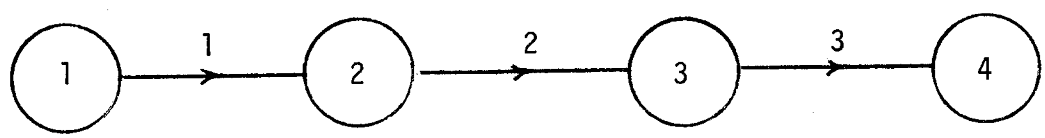

Figure 2 Tandem Link Model: Throughput vs. $\gamma(2)$

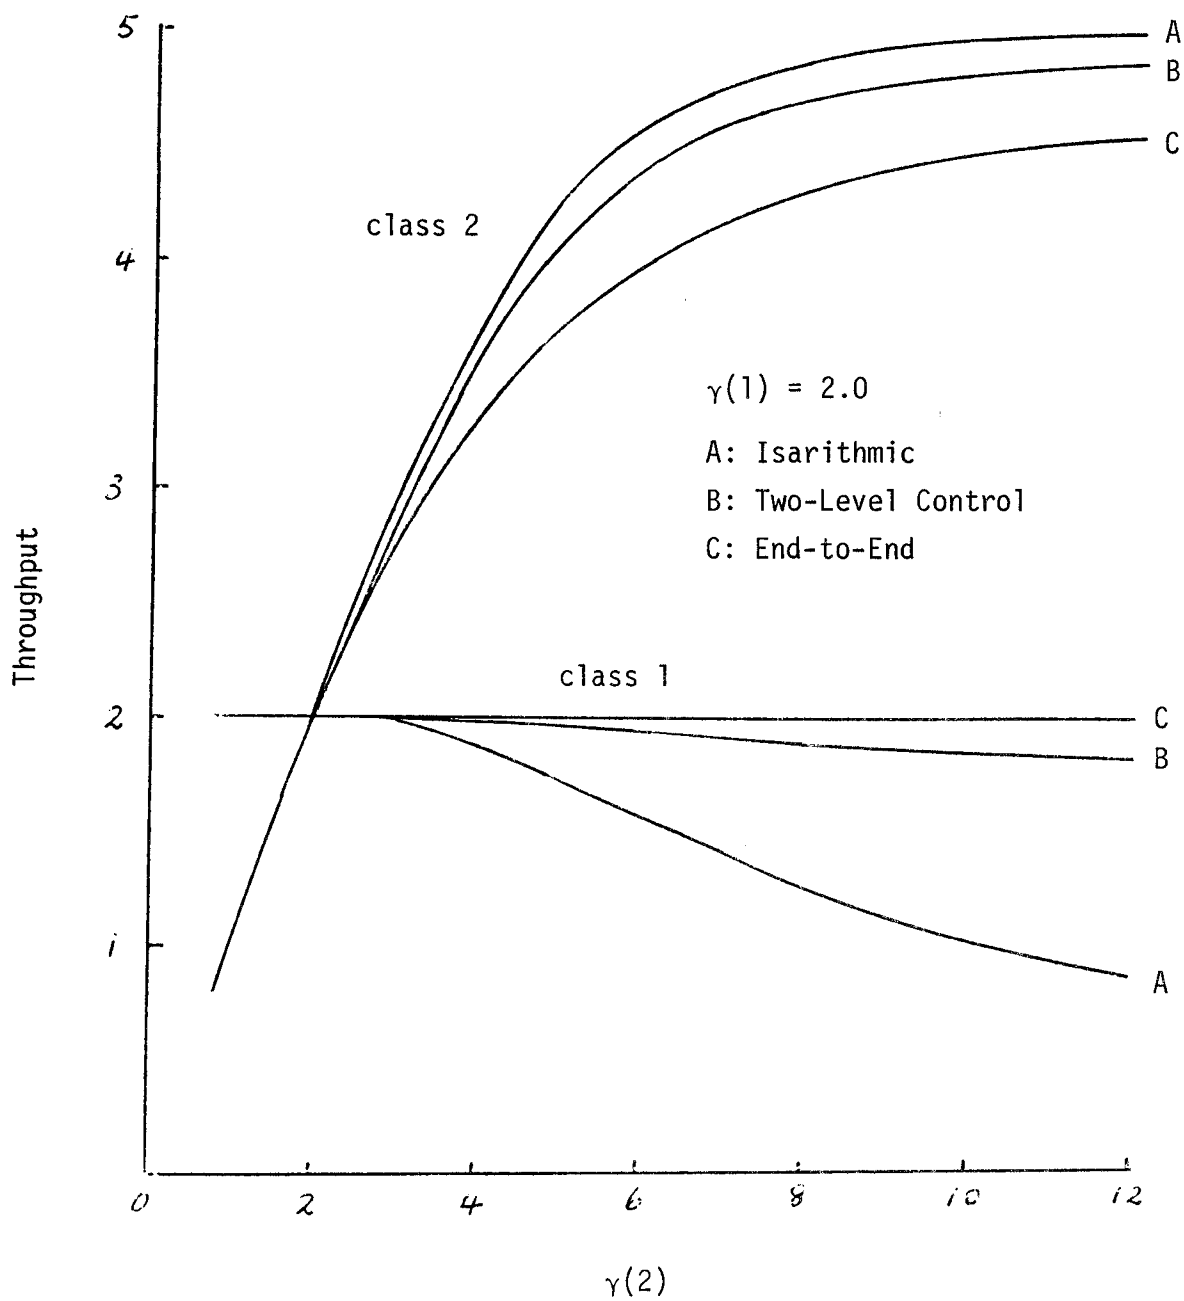

Figure 3 Tandem Link Model: Total Throughput vs. $\gamma(2)$

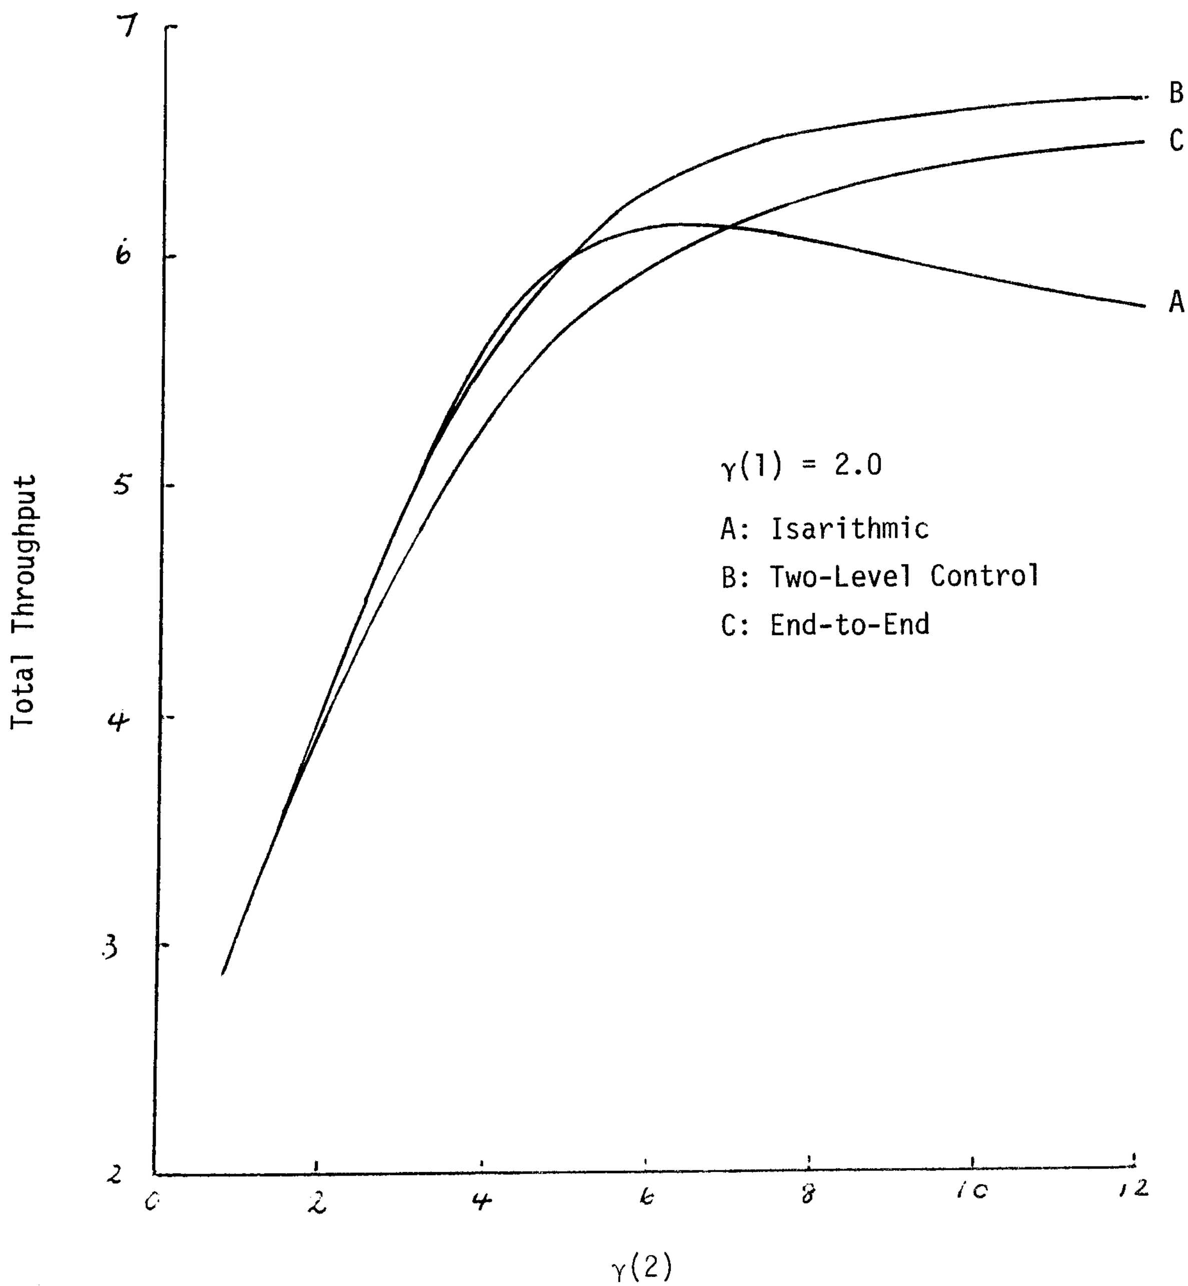

Figure 4 Tandem Link Model: Mean End-to-End Delay vs. $\gamma(2)$

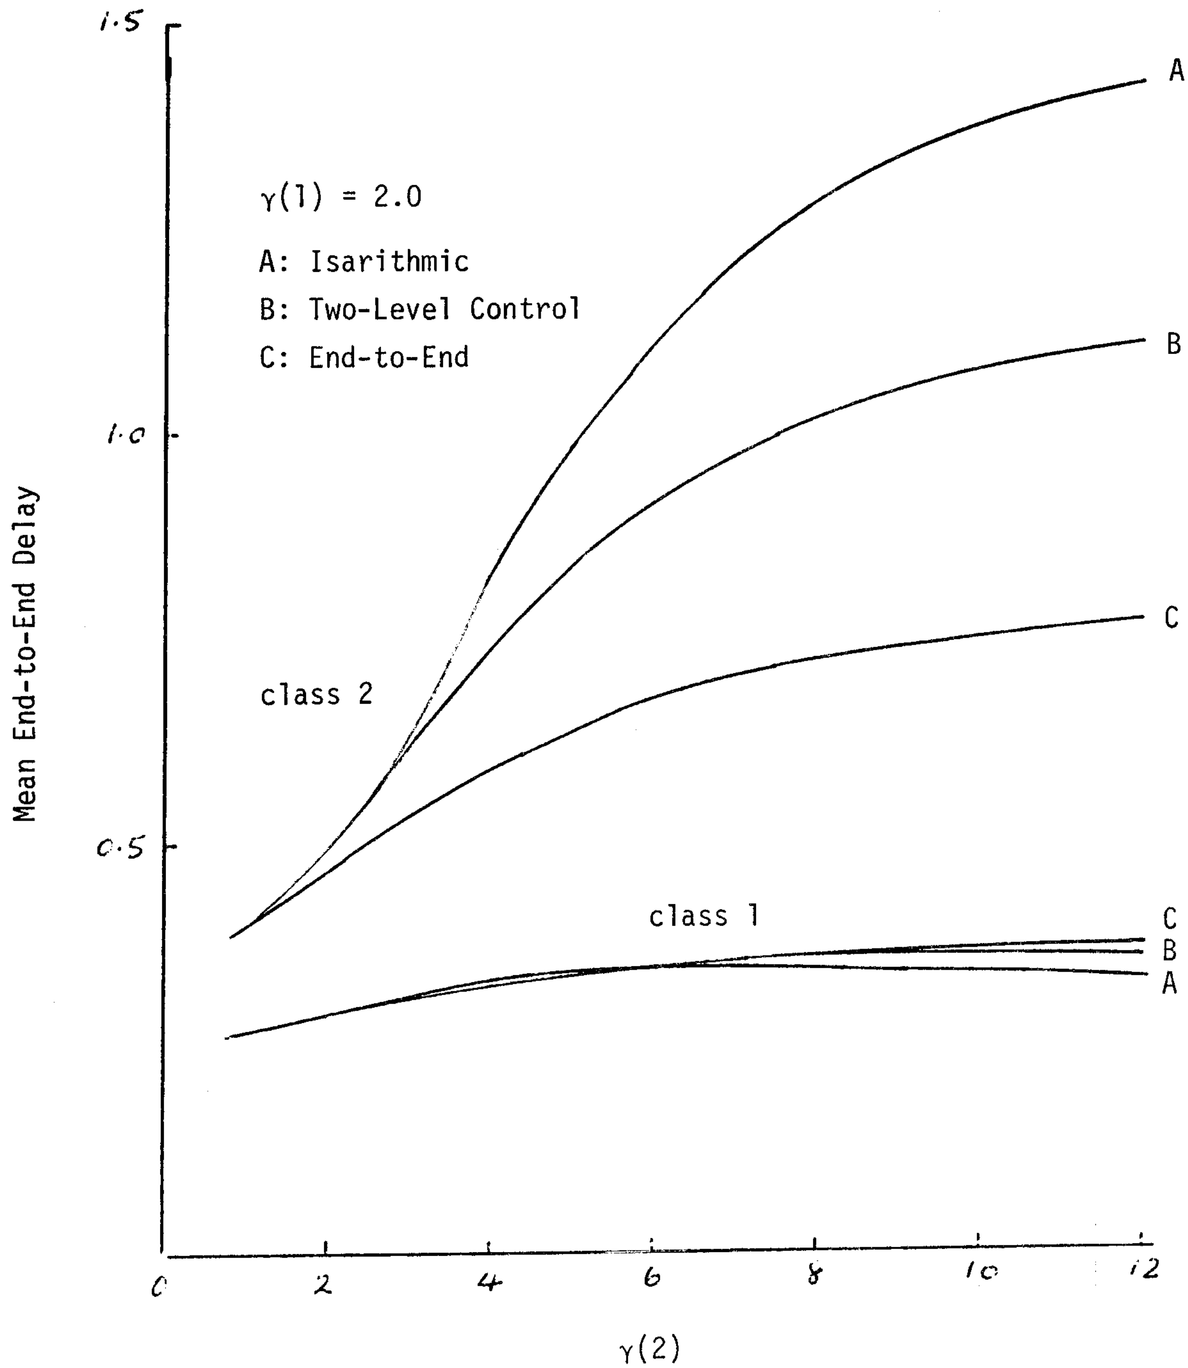

Figure 5 End-to-End Control: Throughput vs. L_2

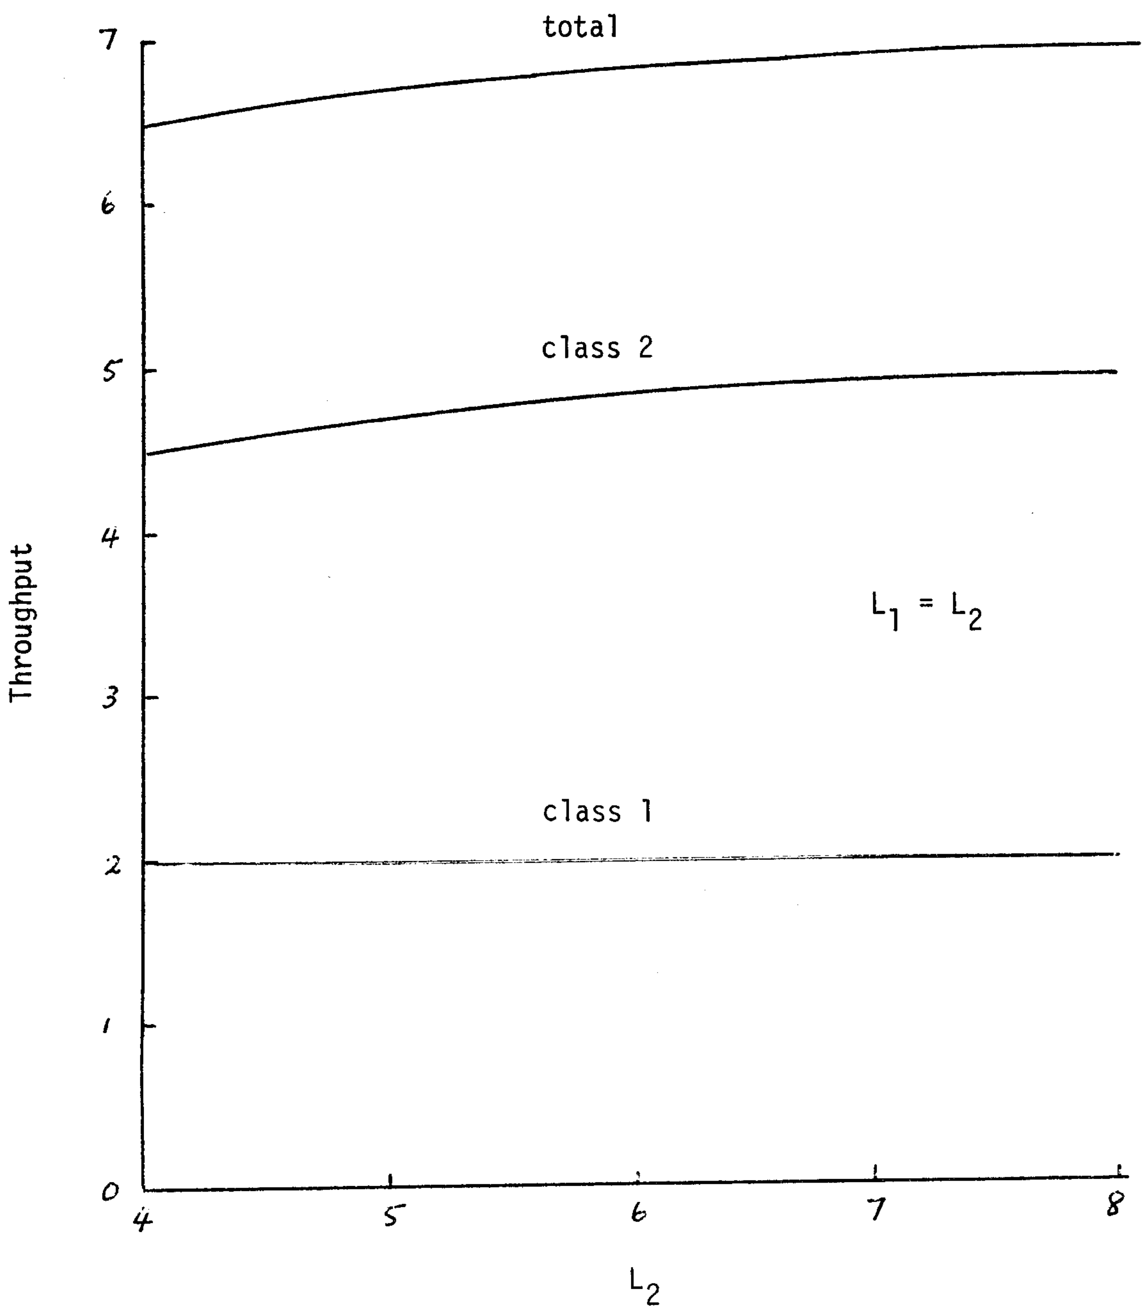

Figure 6 End-to-End Control: Mean End-to-End Delay vs. L_2

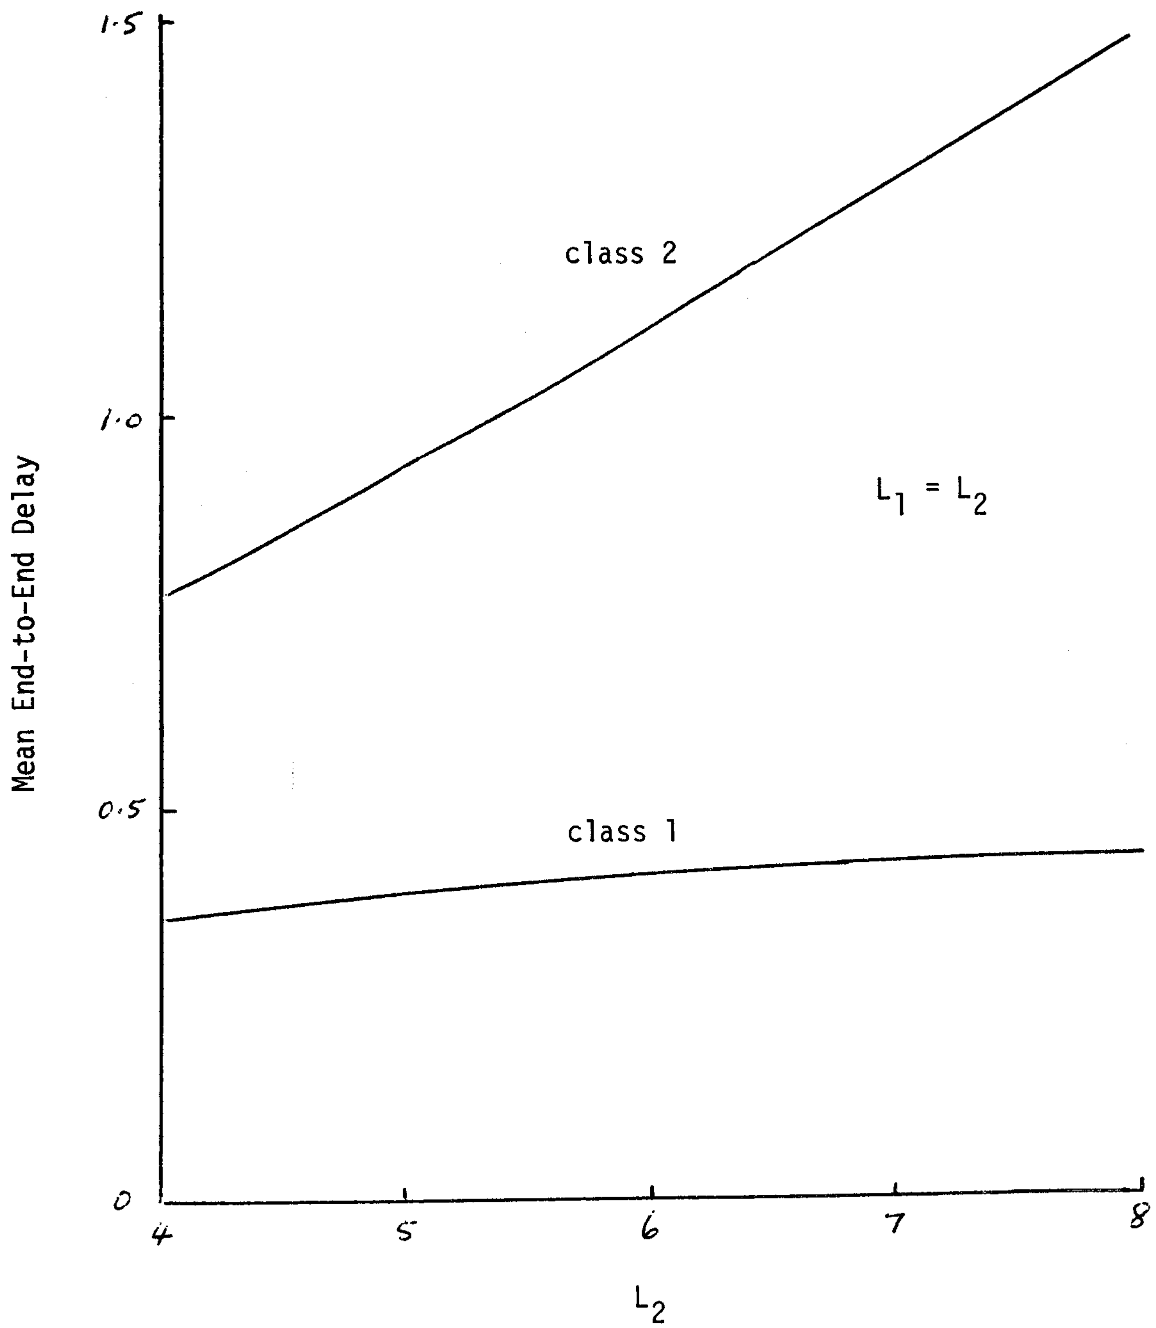

Figure 7 Two-Level Control: Throughput vs. L_2

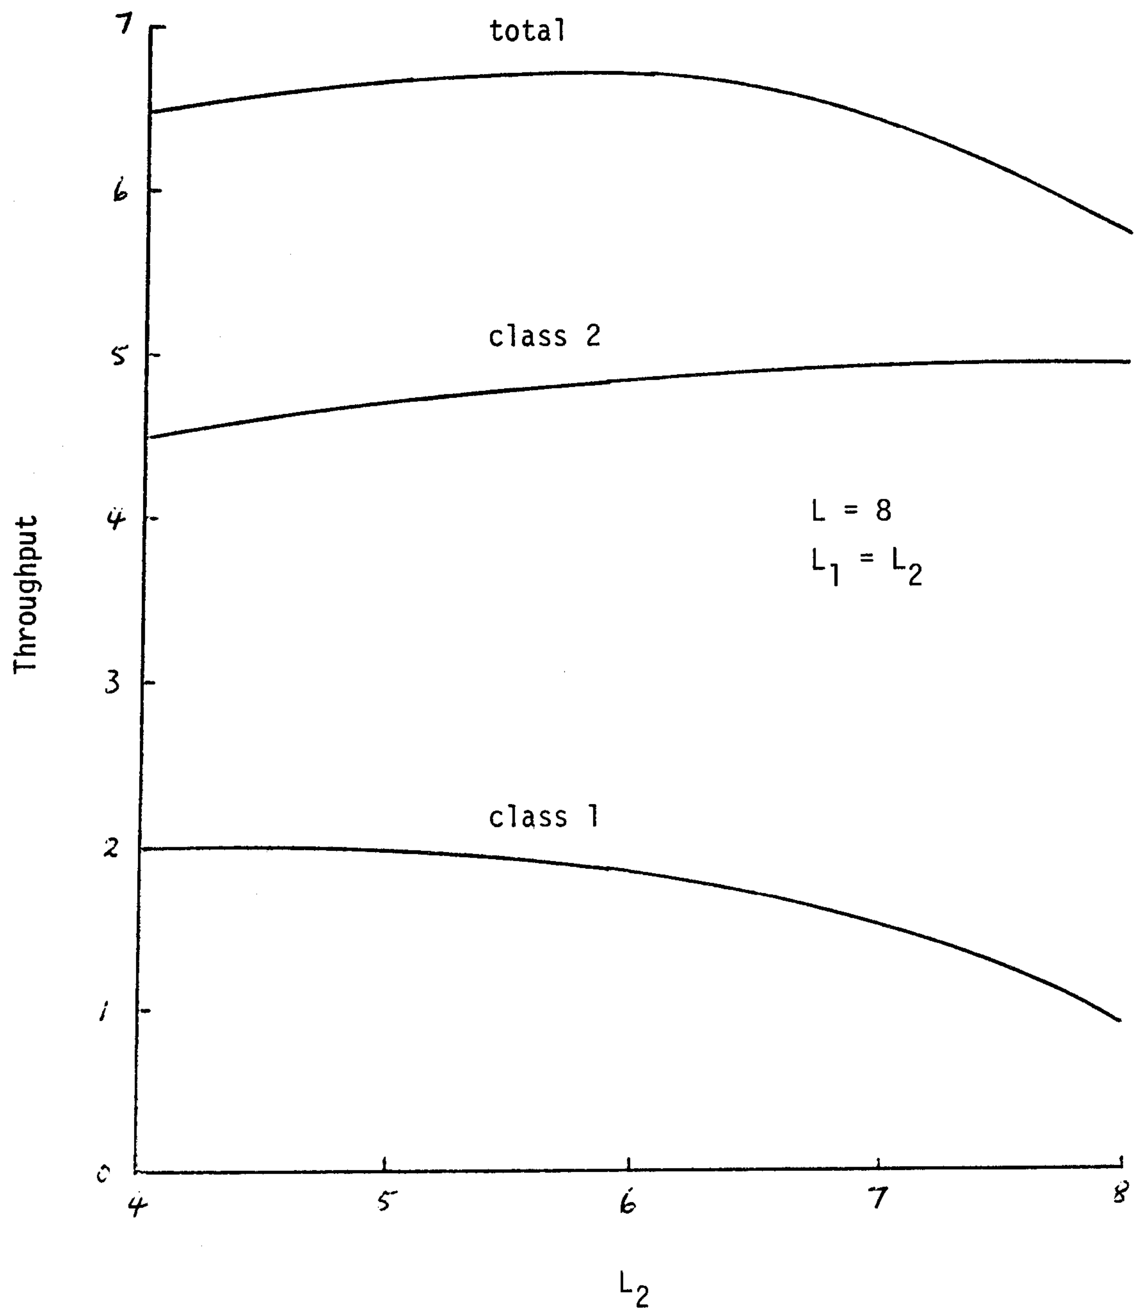

Figure 8 Two-Level Control: Mean End-to-End Delay vs. L_2

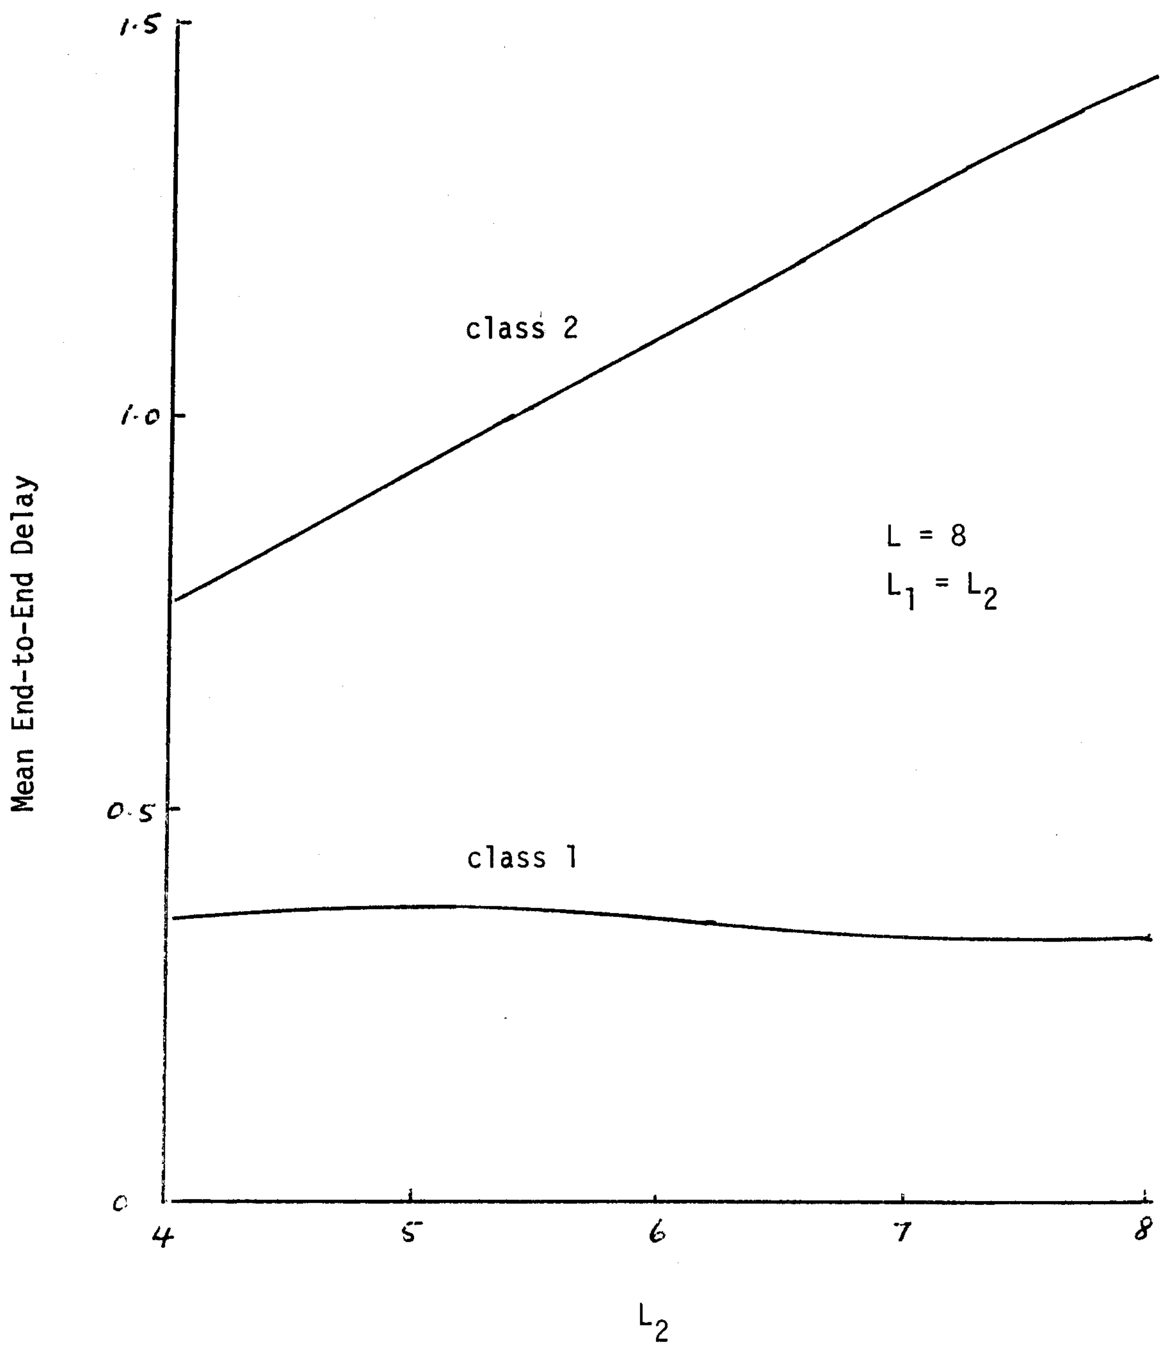

Figure 9 5-node Network

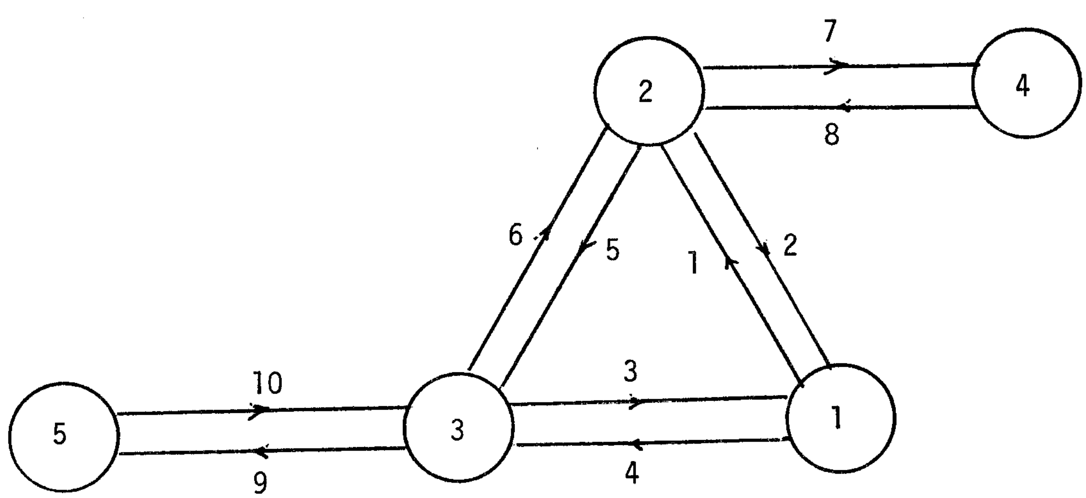

Figure 10 Traffic Matrix

Destination

Source		1	2	3	4	5
	1	0	4α	5α	β	β
	2	4α	0	4α	β	2β
	3	5α	4α	0	2β	β
	4	β	β	2β	0	2β
	5	β	2β	β	2β	0

Figure 11 5-node Network: Throughput of Group 1 vs. β

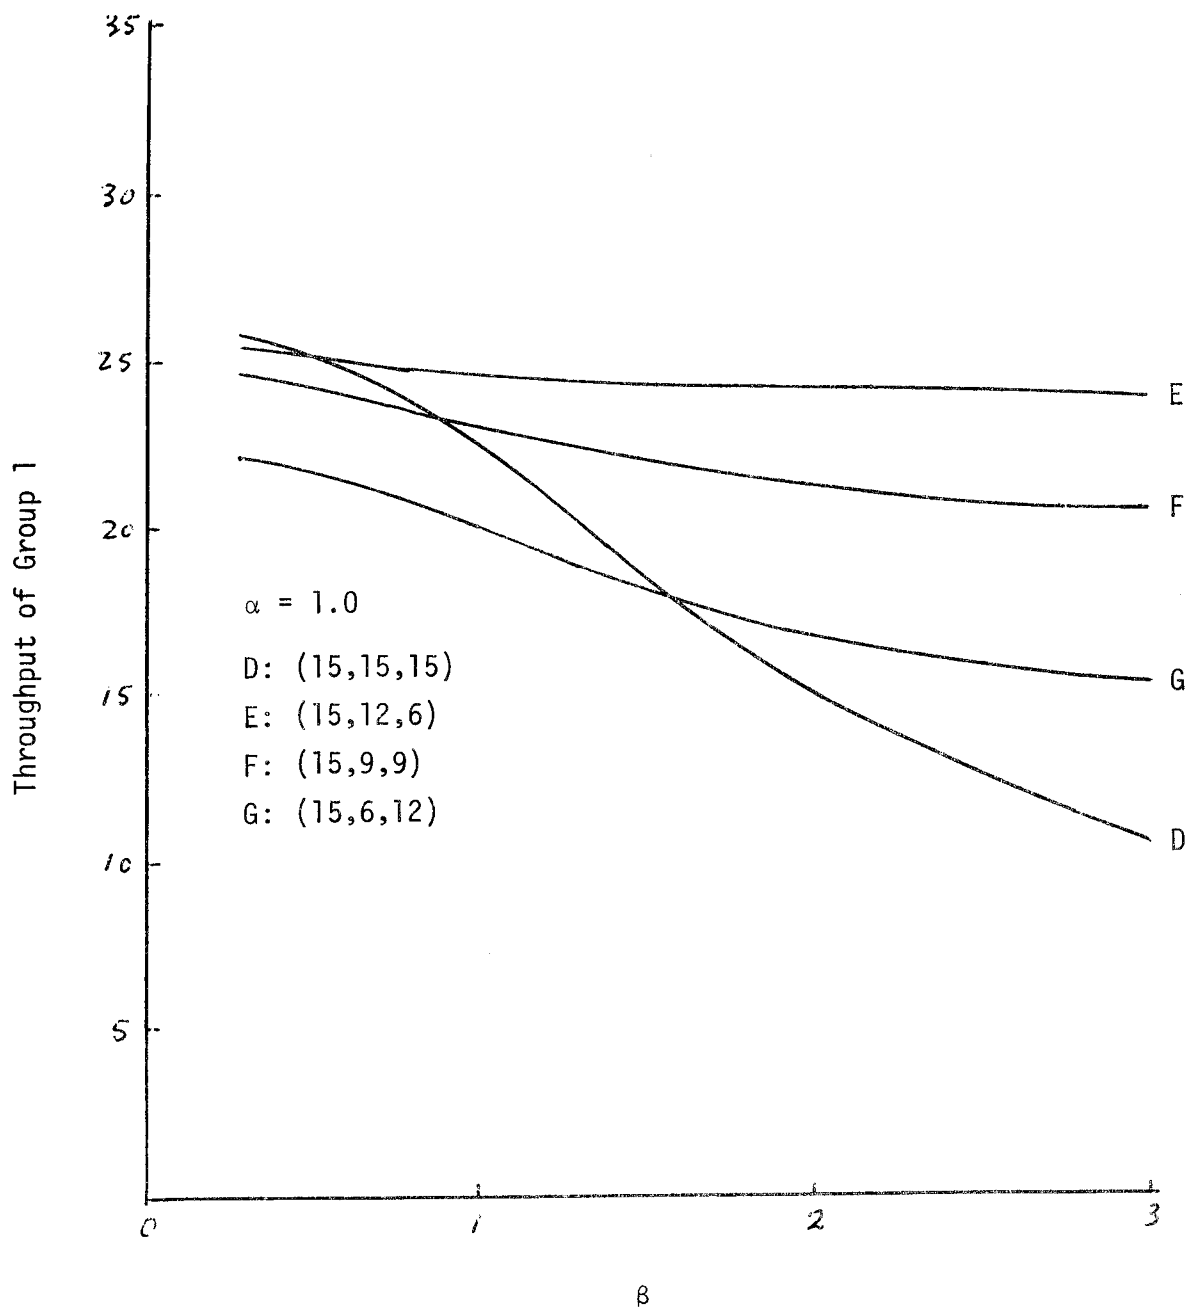

Figure 12 5-node Network: Throughput of Group 2 vs. β

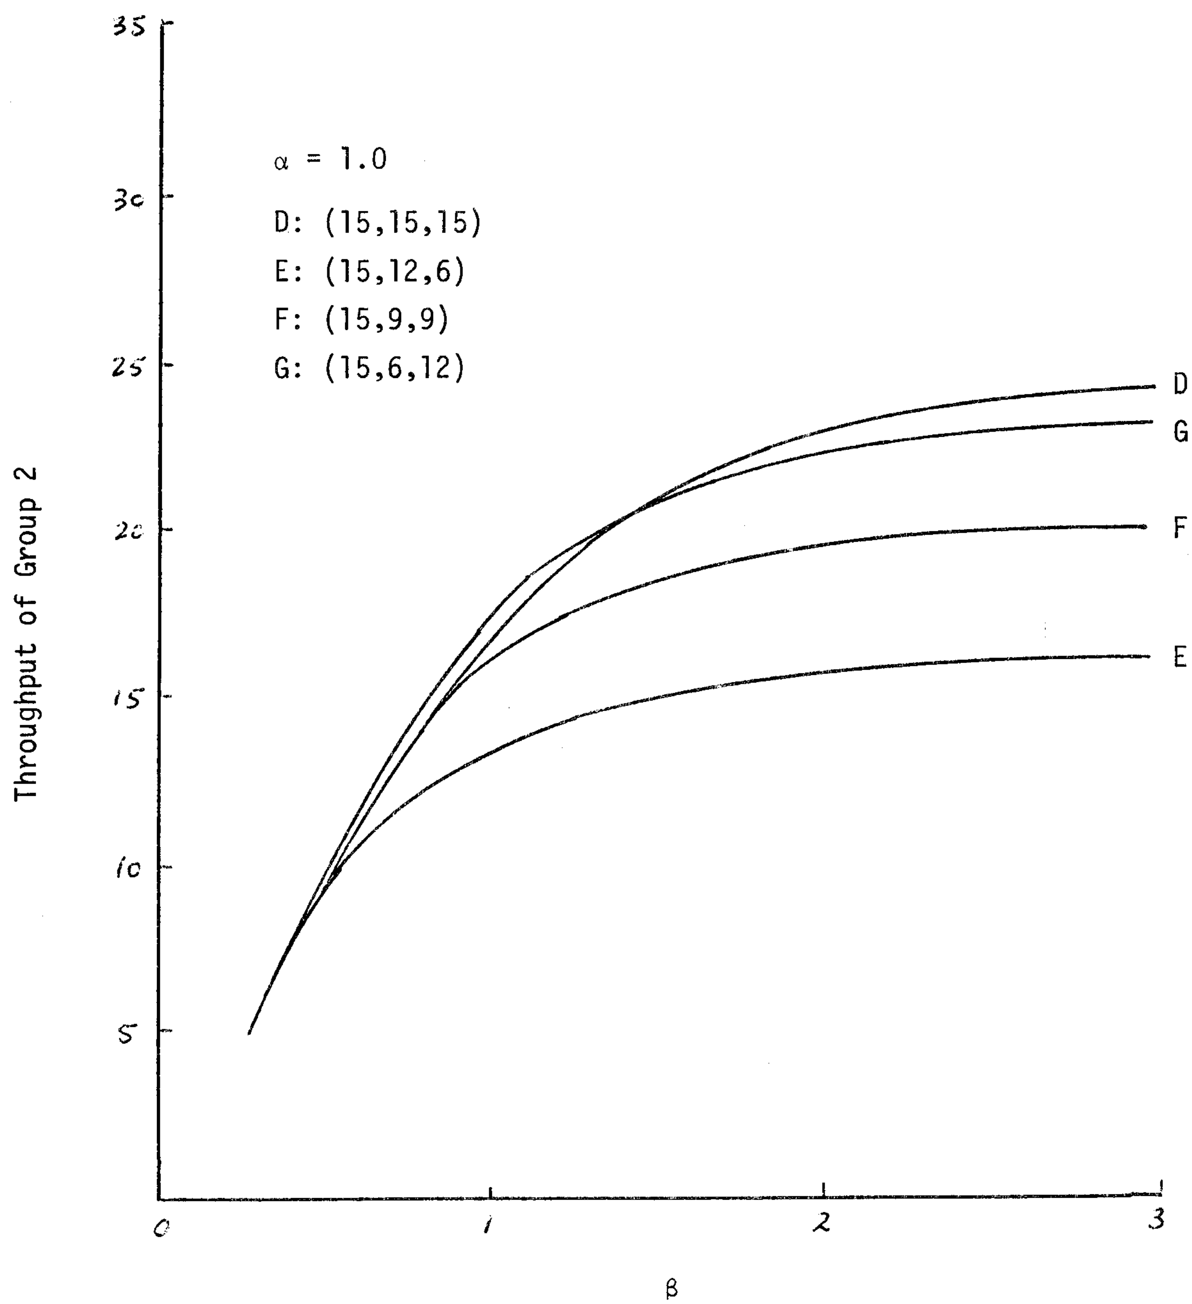

Figure 13 5-node Network: Total Throughput vs. β

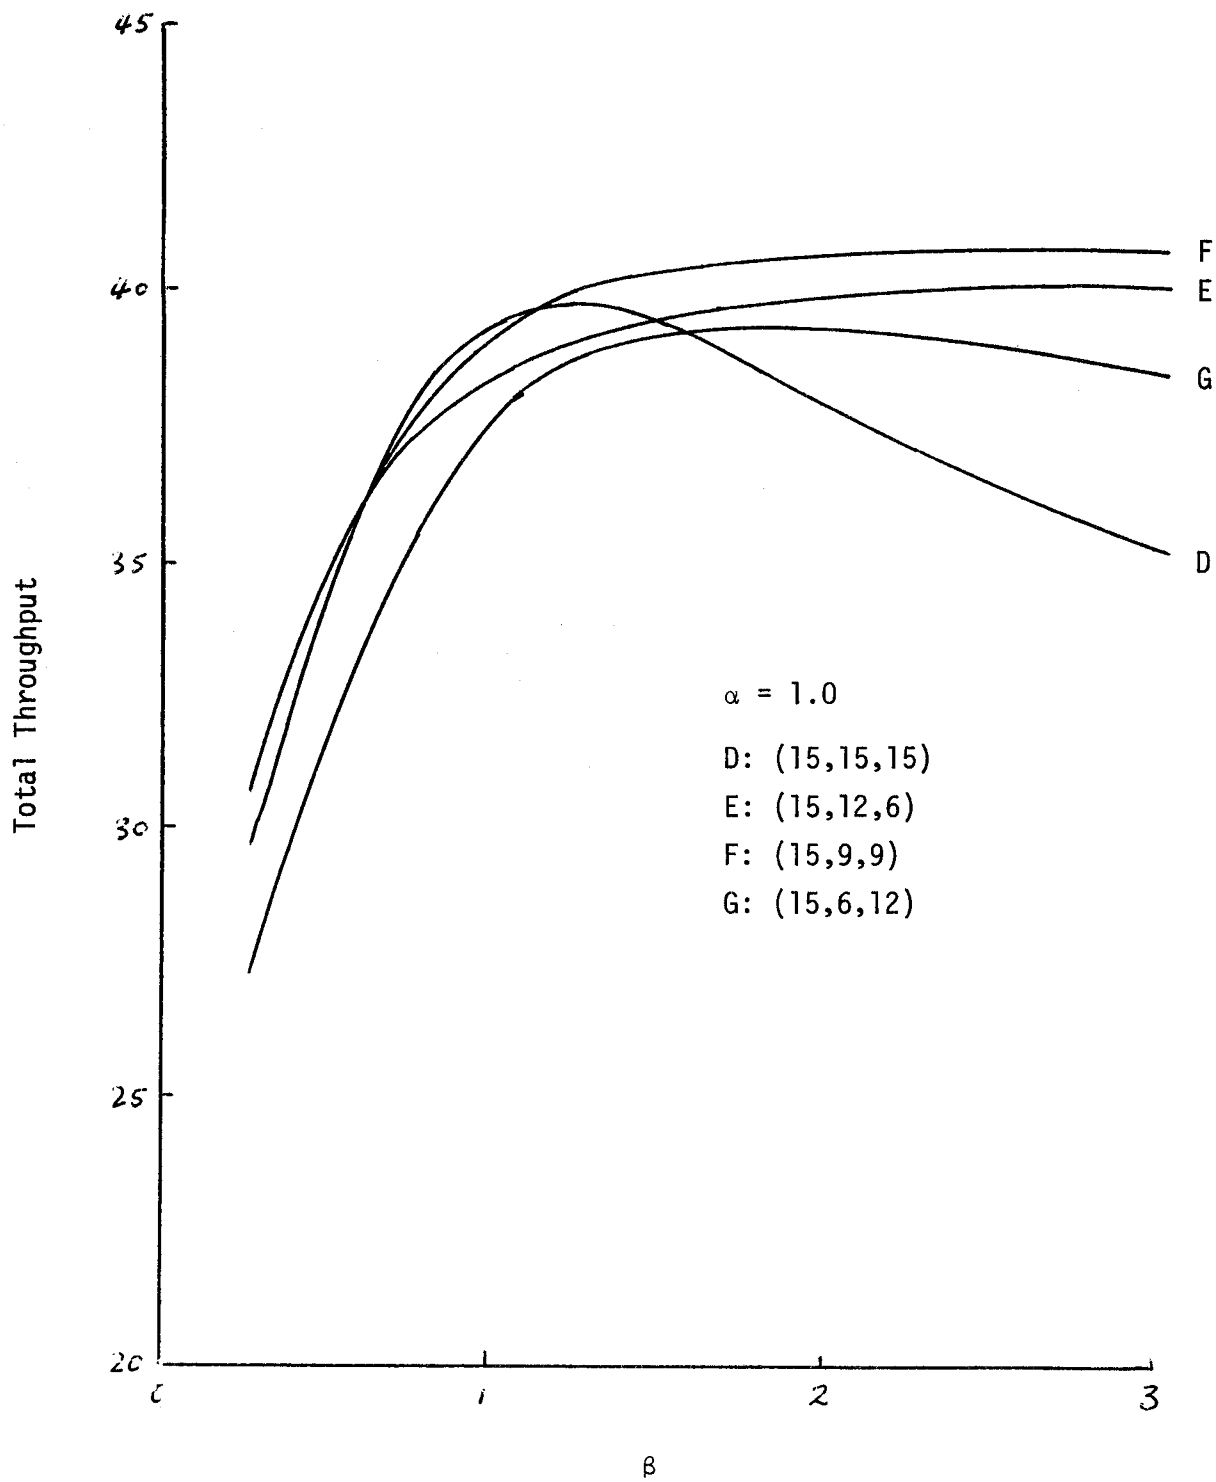

Figure 14 5-node Network: Mean End-to-End Delay vs. β