

**CLASS-DEPENDENT ROUTING IN SWITCHED COMPUTER NETWORKS**

**Irina Neuman**

May 1986

Center for Research on Information Systems  
Information Systems Area  
Graduate School of Business Administration  
New York University

**Working Paper Series**

CRIS #131

GBA #86-81

## 1. Problem definition

A starting point for most of the existing research in the area of backbone network design is the implicit assumption that all messages in the network have similar characteristics and requirements. As a result, a uniform treatment is adopted for all messages, with no distinction being made among different types of applications, each with their own specific characteristics, nor between different user requirements. Though it is true that such an approach greatly reduces the complexity of the analysis, in most cases the assumption does not correspond to the real world environment. Therefore, the solutions generated by algorithms using it, though optimal or near optimal when applied to the simplified problem, may prove to be of poor quality when implemented in practice. Explicitly taking into account the characteristics and requirements of different classes of messages not only will lead to a solution that is preferable from a global perspective (e.g depending on the performance criterion, the average delay in the network may be reduced, or an overall less costly design may be achieved), but also the solution will be better tailored to the individual user needs.

The practical relevance of the issue is also indicated by the fact that routing strategies that differentiate among messages in accordance with their various characteristics and requirements are commonly implemented by many of the existing operational networks. Here are a few examples.

SNA's [Ahuja 79] path control layer defines several classes of service, distinguished by such parameters as response time, security, or reliability requirements. Three transmission priorities (high, medium, and low) are provided. The class of service is specified by the user at session initiation, and is the main criterion on which the choice of the virtual route is based.

DATAPAC [Pandya 77], the public packet-switching network of TransCanada Telephone System, supports two classes of messages, priority and normal. Priorities are user assigned, and are determined by the traffic characteristics. Priority packets are associated with inquiry-response and interactive applications, for which response time is

an important requirement. Normal service, on the other hand, is associated with less time critical applications, such as file transfer and remote job entry. The network has different performance objectives for the two types of flow, namely an average response time of 0.48 seconds for normal packets, and of 0.28 seconds for priority packets. At each node, different routing tables are defined for the two priority classes. At outgoing links, priority packets are given preference (without preemption) over normal packets.

A similar distinction between the two main types of carried traffic is incorporated in the routing strategy of the SITA [Dureste 83, Kroneberg 85] network, a worldwide (and the world's largest) private switched network, that provides telecommunication services to its member airlines. Type A (conversational) traffic consists of very short messages, generally queries/responses for flight reservation, and has very stringent delay requirements (less than three seconds worldwide!). Type B (telegraphic, or conventional) traffic consists of longer messages, of an administrative nature, such as information about flight operations, aircraft movements, lost luggage, etc. The transmission time requirements for these messages are much lower, delivery taking place only overnight in some cases. On the other hand, their reliability requirements are very high, such messages being guaranteed close to 100% security against loss, mutilation, or duplication.

In addition to these real life examples, the pertinence of distinguishing between the priority levels associated with the different classes of messages is also emphasized by the growing body of research literature dealing with the performance analysis of computer networks.

Even when all the characteristics of the network are known, the complexity of the problem is still such that it precludes finding exact solutions for all but the simplest cases. Thus, in [Morris 81] a two node network representing a single full duplex channel is studied. Two classes of messages, with different message sizes and priorities are considered, and the effects on the performance of the simplified network of providing two grades of service between origin and destination are analysed. In addition to the

data messages, the extra traffic generated by acknowledgments is also taken into account. Exact solutions are obtained for three different configurations, distinguished by the types of priorities assigned to messages and to their acknowledgments.

A more general case is studied by Reiser in [Reiser 79], where the analysis is based on approximation methods for closed queuing networks. No distinction is made among the data messages using the network, but control messages, e.g. acknowledgments, are assigned a higher priority. Different message sizes are also considered.

Methods for estimating the average delay, as well as the 90th percentile delay for the two classes of messages supported by the DATAPAC network are developed in [Pandya 77]. Message behavior at a link in the network is modelled as a single server, head-of-the-line priority system with two priority classes. Exact and approximate models are developed and used to analyse the sensitivity of message delays to changes in various system parameters, such as link capacity and utilization, message length, and the relative mix of priority and normal messages.

In [Laue 81] the effect on network delay of such parameters as message versus packet switching, type of traffic carried by the network, or the presence of a priority discipline at the queues is studied. The data provided by the Bell System's OSN network [Amoss 80] is analysed and interpreted. The main conclusion of the study is that for a network that, in addition to short inquiry-response traffic, also carries longer messages, the use of a priority discipline is essential in achieving a good level of service. Moreover, the numerical results indicate that the introduction of longer messages has little effect on the delay incurred by the high priority traffic.

The comparative results in the area of performance evaluation of computer communication systems supporting several classes of service strongly suggest that the overall performance is significantly improved when messages are prioritized. These theoretical indications, together with the experience gained from the networks that chose to implement similar methods, are powerful arguments in favor of such schemes. Nevertheless, the literature dealing with the related design issues is very limited, the only relevant work being a series of papers by K. Maruyama and D. T. Tang.

In [Maruyama 76], messages are classified according to their processing and delay characteristics, and a priority level is associated with each message class. The procedure is further refined in [Maruyama 77] and [Tang 76], where the priority levels are no longer assumed to be known in advance, and as a result a better overall design may be achieved. In all cases, simple heuristic procedures are used for solving the models.

The problems considered here, as well as other related network design problems, have an extremely complex structure. This fully justifies the use of heuristic methods for their solution. But an important shortcoming of much of the existing work in the field is that no means, theoretical or empirical, is available for evaluating the quality of the solutions provided by the heuristic, which may seriously hamper their usefulness for real life applications. In addition to representing a new perspective into some important aspects of the problem, the approach presented here has the advantage of also providing, as part of the suggested algorithm, for a benchmark against which the quality of the generated solution is tested.

## 2. Problem formulation

We consider a network with a given topology, where the link capacity values are also known. Since we deal with a design problem, the amount of flow the network will have to carry cannot be known in advance, instead traffic estimates have to be used. We assume that messages from each class arrive on the boundaries of the network according to Poisson processes with given average interarrival times, and that message lengths are exponentially distributed for each class. Further details about the modeling of the queuing phenomena in the network can be found in [Neuman 86].

A message arriving at a network switch is placed in a waiting queue, until the line on which it must be forwarded next becomes available. Each message class is associated with a known priority level. A head-of-the-line non-preemptive discipline is imposed on the messages waiting for each link. In this external priority scheme, messages queue according to the priority class they belong to, i.e a message arriving at a node joins the queue behind all messages of equal or higher priority, but ahead of those of lower priority groups already waiting.

For each communicating pair of network nodes, the designer must provide the model with a set of candidate routes, out of which the choice of primary routes to be used for each traffic class will be made. After an initial solution is generated, the designer could continue by interactively changing the structure of the candidate sets, and comparing the resulting solutions.

The following notation will be used throughout the paper:

$L$	=total number of links in the network
$J$	=total number of priority classes
$1/\mu_j$	=average message length for class $j \in J$
$Q_l$	=capacity [bps] of link $l \in L$
$D_j$	=unit cost of delay for messages in class $j \in J$
$R$	=set of candidate routes
$\Pi$	=set of communicating origin-destination pairs in the network
$S_p^j$	=set of candidate routes for class $j$ messages associated with origin-destination pair $p \in \Pi$ .
	$S_p, p \in \Pi$ is defined as $\cup_{j \in J} S_p^j$ .
$\lambda_r^j$	=the class $j$ message arrival rate for the unique origin-destination pair associated with route $r \in R$ . Also, $\lambda_p^j = \lambda_r^j, \forall r \in S_p^j$
$\phi_l^j$	=the class $j$ message rate on link $l$
$F_l^j = \phi_l^j / \mu_j$	=the class $j$ bit flow on link $l$
$T_l^j$	=the average delay incurred on link $l$ by a class $j$ message.
$x_r^j$	=a decision variable, taking the value 1 if route $r$ is chosen to carry the class $j$ flow of its associated origin-destination pair, and 0 otherwise.

The choice of routes plays a very important role in determining the delay experienced by messages travelling through the network. As a result, we chose our performance criterion to be the total queuing cost, i.e the total cost of the delay associated with all message classes in the network over a given time period. Computing this measure hinges on being able to express the time that a message is delayed, on an average, at each of the links in the network. Thus  $T_l^j$ , the average delay on link  $l$  for class  $j$  messages is ([Kleinrock 76]):

$$T_l^j = \frac{s_l^j(1-\sigma_j) + \sum_{i=j}^J \phi_l^i (s_l^i)^2}{(1-\sigma_j)(1-\sigma_{j+1})} \quad (1)$$

where  $s_l^j = 1/\mu_j Q_l$  is the average transmission time for class  $j$  messages on link  $l$ , and  $\sigma_j = \sum_{i=j}^J \phi_l^i s_l^i$ .  $T_l^j$  includes both the queuing delay incurred by a message while waiting in the buffers of a network switch, as well as the transmission time.

The above expression becomes untractably complex as the number of message classes increases. In the following we will therefore concentrate on the case of a network supporting just two classes of messages. It is an important and not very restrictive case, as indicated for instance by the sharp distinction between traffic generated by interactive computation, with its tight delay requirements, on one hand, and such applications as file transfer and remote job entry, for which response time is less of a critical factor, and which as a result may be associated with a lower priority, on the other. The relevance to real life applications of such a dichotomy is also attested by the implementation examples mentioned earlier, such as the DATAPAC and the SITA networks.

We will therefore think of all the messages in the network as being grouped into two classes, and without loss of generality, we will assume that the higher priority is associated with the second class.

From 1, the following expressions are obtained for the average delay on link  $l$  for class 1 and class 2 messages, respectively:



$$\begin{aligned}
T_l^1 &= \frac{(Q_l - F_l^2)/\mu_1 + F_l^2/\mu_2}{(Q_l - F_l^1 - F_l^2)(Q_l - F_l^2)} \\
T_l^2 &= \frac{1}{\mu_2(Q_l - F_l^2)}
\end{aligned} \tag{2}$$

The average class  $j$  bit flow on link  $l$  can be stated in terms of the decision variables  $x_r^j$  as:

$$F_l^j = \sum_{r \in R} \lambda_r^j \delta_{rl} x_r^j / \mu_j \tag{3}$$

It is then possible to express a general cost function that reflects how the two classes of messages are affected by the delay they experience in the network:

$$C = \sum_{j \in J} D_j \sum_{l \in L} \phi_l^j T_l^j = D_1 \sum_{l \in L} \frac{F_l^1(Q_l - F_l^2) + a F_l^1 F_l^2}{(Q_l - F_l^1 - F_l^2)(Q_l - F_l^2)} + D_2 \sum_{l \in L} \frac{F_l^2}{Q_l - F_l^2}$$

where  $a = \mu_1/\mu_2$ , and the  $F_l^1$  and  $F_l^2$  are defined by 3 and are used here for simplicity of notation.

The nature of the problem imposes certain restrictions upon the characteristics of the higher priority messages. Their length cannot exceed a certain limit, and they have to pay for the increase in performance they require. As a result, the following relations must hold among the problem parameters:

1.  $a \leq 1$ , i.e the average length of class 2 messages cannot exceed that of class 1 messages, and
2.  $D_2 \geq D_1$ , i.e the unit cost of delay is at least as high for class 2 messages as for class 1 messages.

The structure of the problem becomes more apparent if a new set of decision variables is introduced. The utilization of each link  $l$ , i.e that portion of its capacity that is actually used, can be expressed as:



$$f_l = \frac{F_l^1 + F_l^2}{Q_l} = f_l^1 + f_l^2$$

where  $f_l^j$ ,  $j=1,2$ , is that portion of the link utilization attributable to class  $j$  flow. Far from increasing the complexity of the problem, the introduction of this set of derived decision variables better highlights its underlying structure, and results in a more effective solution procedure.

In this context, the problem of optimally assigning primary routes in a network supporting two classes of messages becomes that of selecting the  $x_r^j$  and  $f_l^j$  values that satisfy:

### Problem P1

$$Z_{P1} = \min \left\{ D_1 \sum_{l \in L} \frac{f_l^1(1-f_l^2) + a f_l^1 f_l^2}{(1-f_l^2)(1-f_l^1-f_l^2)} + D_2 \sum_{l \in L} \frac{f_l^2}{1-f_l^2} \right\} \quad (4)$$

subject to:

$$\sum_{r \in R} \lambda_r^j \delta_{rl} x_r^j / \mu_j Q_l \leq f_l^j \quad \forall l \in L, j=1,2 \quad (5)$$

$$f_l^1 + f_l^2 \leq 1 \quad \forall l \in L \quad (6)$$

$$\sum_{r \in S_p^j} x_r^j = 1 \quad \forall p \in \Pi, j=1,2 \quad (7)$$

$$f_l^j \geq 0 \quad \forall l \in L, j=1,2 \quad (8)$$

$$x_r^j = 0,1 \quad r \in R, j=1,2 \quad (9)$$

The constraint set ensures that the total flow on each link is feasible in terms of the capacity value of the link, and that only one route is chosen for each message class associated with a given origin-destination pair.

The formulation of the problem implicitly takes into account the different delay

requirements of the two classes of messages. Priority messages, with their tighter response time requirement, are associated with a higher cost of delay, which will lower even further the average delay they incur in the final solution. In other words, messages that demand a better service have to pay for it, and the quality of the service reflects the price paid. This flexibility is introduced in the model by the fact that, as an alternative to the traditional minimization of the average message delay in the network, the optimization criterion used is the minimization of an estimate of the overall queuing cost, expressed as a function of the specific unit costs of delay associated with each message class.

### 3. Solution procedure

It can be shown that the problem defined in the previous section belongs to the NP-complete class (see [Neuman 86]). As a result, it is extremely hard to solve to optimality. Fortunately, since we only deal with traffic and cost estimates, the very nature of the problem renders an exact solution unnecessary. We will therefore outline in the following ways to obtain good feasible solutions (that also represent upper bounds on the value of the optimal solution). We will also show how lower bounds on the optimal value can be generated. Since the value of the optimal solution lies somewhere between the best upper and lower bounds obtained, this bounding technique provides for an effective way to ascertain the quality of the heuristic solution.

Here, the upper and lower bounding procedure will only be summarized. Further mathematical and algorithmical details can be found in [Neuman 86].

The lower bound is obtained by relaxing the constraints in 5, and a Lagrangean problem is formed by adding them into the objective function, together with appropriate multipliers. With the coupling constraints no longer present, the problem becomes decomposable into  $|L|+2\times|H|$  subproblems, one for each link, and for each origin-destination pair and message class. While the origin-destination subproblems are very easy to solve, the structure of the link subproblems is more intricate, and numerical methods are required for their solution.

The lower bound thus obtained is improved upon by applying a subgradient optimization procedure, an iterative technique that attempts to close the gap between the Lagrangean and optimal values. This method was found to be effective in a variety of combinatorial problems (see for example [Held 71, Gavish 82, Gavish 83, Gavish 86]), and performed well also in this case.

A search for feasible solutions is also carried out at each subgradient iteration, using the solution to the Lagrangean problem (not necessarily feasible!) as a starting point. In addition to testing the Lagrangean solution for feasibility, a simple but effective procedure was also incorporated in the system. At each iteration, and for each origin-destination pair, a list of "good" candidates for primary routes is generated. Then, a route is randomly chosen from each list, and each alternative solution is in turn checked for feasibility. Whenever a feasible solution of lower cost is generated, its value and associated route choices become the best current solution, and the search continues.

#### **4. Computational Results**

The model and solution procedures presented earlier are currently implemented in a system that allows for an easy and flexible definition of the network characteristics and model parameters. At the end of each major iteration (defined as a number of subgradient iterations, to be specified by the user), control is returned, and the procedure can be either stopped, if an acceptable solution was reached, or continued. Also, a comprehensive output, depicting the details of the best feasible solution generated so far can be viewed by the user at this point.

Four network topologies were used for testing the model. They appear in figures 1 through 4. The total average message traffic for all origin-destination pairs is of four messages for both directions, evenly divided between the two types of flow.

The performance of the algorithms, as well as the way in which various changes in the parameter values are reflected in the general characteristics of the solution generated, can be observed in the tables 1 through 3.

For the results in table 1, the average length for class 2 messages is fixed to 200 bits, while that of class 1 messages is allowed to vary from 400 to 800 bits. The unit costs of delay used for all problems in this table are  $D_1=1000$  and  $D_2=4000$ . For low to medium loads, the algorithm performs extremely well. For these cases, the gap between the lower and the upper bounds is so small, that it practically ensures that an optimal, or very close to optimal, solution is reached. For higher loads, the performance somewhat worsens, but the gaps are still reasonably small, so that the generation of good feasible solutions is guaranteed. In all cases, the convergence was fast, the results being obtained in 50 to 100 subgradient iterations.

The link capacities are such that, in all cases, for  $1/\mu_1=400$  bits the networks are very lightly loaded, while for  $1/\mu_1=800$  bits many links operate close to saturation. In spite of this, the average delay for priority messages is practically unchanged, while the delay for class 1 messages increases drastically, i.e the quality of service that class 2 messages receive is only marginally affected by changes in the volume of nonpriority traffic. It is known that, at the link level, the behavior of priority messages is not influenced by the presence of nonpriority traffic (see 2). The numerical results suggest that, in spite of the more complex interactions present in this case, this also holds true at the network level. On the other hand, the average delay for class 1 messages increases significantly, and the degradation in performance is proportionally higher as the message length increases and some of the links in the network approach saturation. For instance, in the case of the RING network, an increase from 600 to 700 bits per message results in a 38.7% increase in the delay, while a further increase in message length to 800 bits corresponds to a jump of 53.7% in delay.

Table 1 also contains the solutions obtained for the same parameter values, under a first-come first-served discipline, i.e the case when no distinction is made between the two types of traffic. The results are needed for comparative purposes, to determine to what extent the overall performance is improved as a result of using a prioritizing scheme.

Notice that in all but one case, as a consequence of no longer giving preferential treatment to shorter messages, the overall cost obtained for the FCFS case is higher, the difference being quite significant for some cases. The distribution of the cost between the two classes changes: higher than in the priority case for class 2 messages, and correspondingly lower for class 1. The only exception are the results for the GTE network with  $1/\mu_1=800$  bits, when both costs increase. This could be explained by the relatively large gap between the lower and the upper bounds, and it is expected that the same characteristic would not be present in the optimal solution. The changes in cost structure are reflected in the new average delays for the two classes of messages. Even for the one case when the FCFS scheme generates better results (the GTE network with  $1/\mu_1=600$  bits) the delay for class 1 decreases, while the delay for class 2 went up, but the difference was too small to be reflected in the overall cost. An extreme example of degradation in performance is the case of the RING network with  $1/\mu_1=800$  bits. Under this heavy load, the impact of using a priority scheme is even more marked. Thus, for the FCFS case, even after a large number of iterations, the best solution generated still assigns a very high delay to class 2 messages. Also notice that, as a result of no longer receiving preferential treatment at the links in the network, the average delay for class 2 messages becomes sensitive to changes in the overall load, and increases as the average message length for class 1 goes up.

Another set of experiments is aimed at analysing the effects on the performance of the network of keeping a fixed traffic load, while changing the ratios between the two types of traffic. The total amount of external traffic generated by each origin-destination pair (expressed in bits) is held constant, and the average message lengths for the two message classes are varied accordingly. The results are summarized in table 2. The first problem, for each network, corresponds to the equal message length case. Thus the difference in performance is solely explained by the queuing discipline used at the links in the network, and the results clearly evidence the impact of using a priority scheme. As expected, when the weight in the total traffic of priority flow is reduced, their associated queuing cost decreases, while the queuing cost for class 1 flow goes up. But, due to their lower unit cost of delay, the effect of this increase is more than

compensated by the improved level of service class 2 messages receive, and the overall cost is reduced. The only exception from this general trend is shown by the RING network, where the queuing cost for class 1 messages is highest for the equal message length case. This is reflected in the fact that, in this case, the corresponding difference between the average message delays for the 2 classes is more significant (67%, as opposed to 26% to 55% for all other cases), i.e the effects of the priority scheme are even more marked.

For the problems summarized in table 3 the unit cost of delay for priority messages, is hold costant at 4000, while  $D_1$  is allowed to vary. The changes in the ratio between the two costs of delay does not seem to influence the delays experienced by priority messages (only for the  $D_1=D_2$  case, a small increase in the delay for class 2 messages is observed). The performance of nonpriority messages, on the other hand, is affected; only slightly while the two costs of delay are commensurable, more significantly as  $D_1$  goes to zero. The fact that this degradation is not paralleled by an improved performance for priority messages, it suggests that the results are due to the decreased weight in the objective function of the component corresponding to class 1 messages, i.e their delay becomes unimportant. This fact is also reflected in the larger gap between the lower and the upper bounds obtained for these cases.

## 5. Summary and Conclusions

A model and solution method for the problem of assigning primary routes to each of the classes of messages carried by a network with known topology and link capacities, was presented. The modelling approach implicitly takes into consideration the different requirements associated with the two types of traffic. Based on extensive computational experiments, it can be concluded that the convergence of the algorithm is good, and feasible solutions that are very close to optimality are generated. Moreover, the different types of tests performed, have evidenced the impact that distinguishing between the two types of flow has on the general characteristics of the solution reached, thus stressing the importance of taking these differences into account in the design phase.







Figure 1: Topology and link capacities (Kbps) for the GTE network





**ARPA**  
**21 NODES**  
**26 LINKS**

Figure 2: Topology and link capacities (Kbps) for the ARPA network

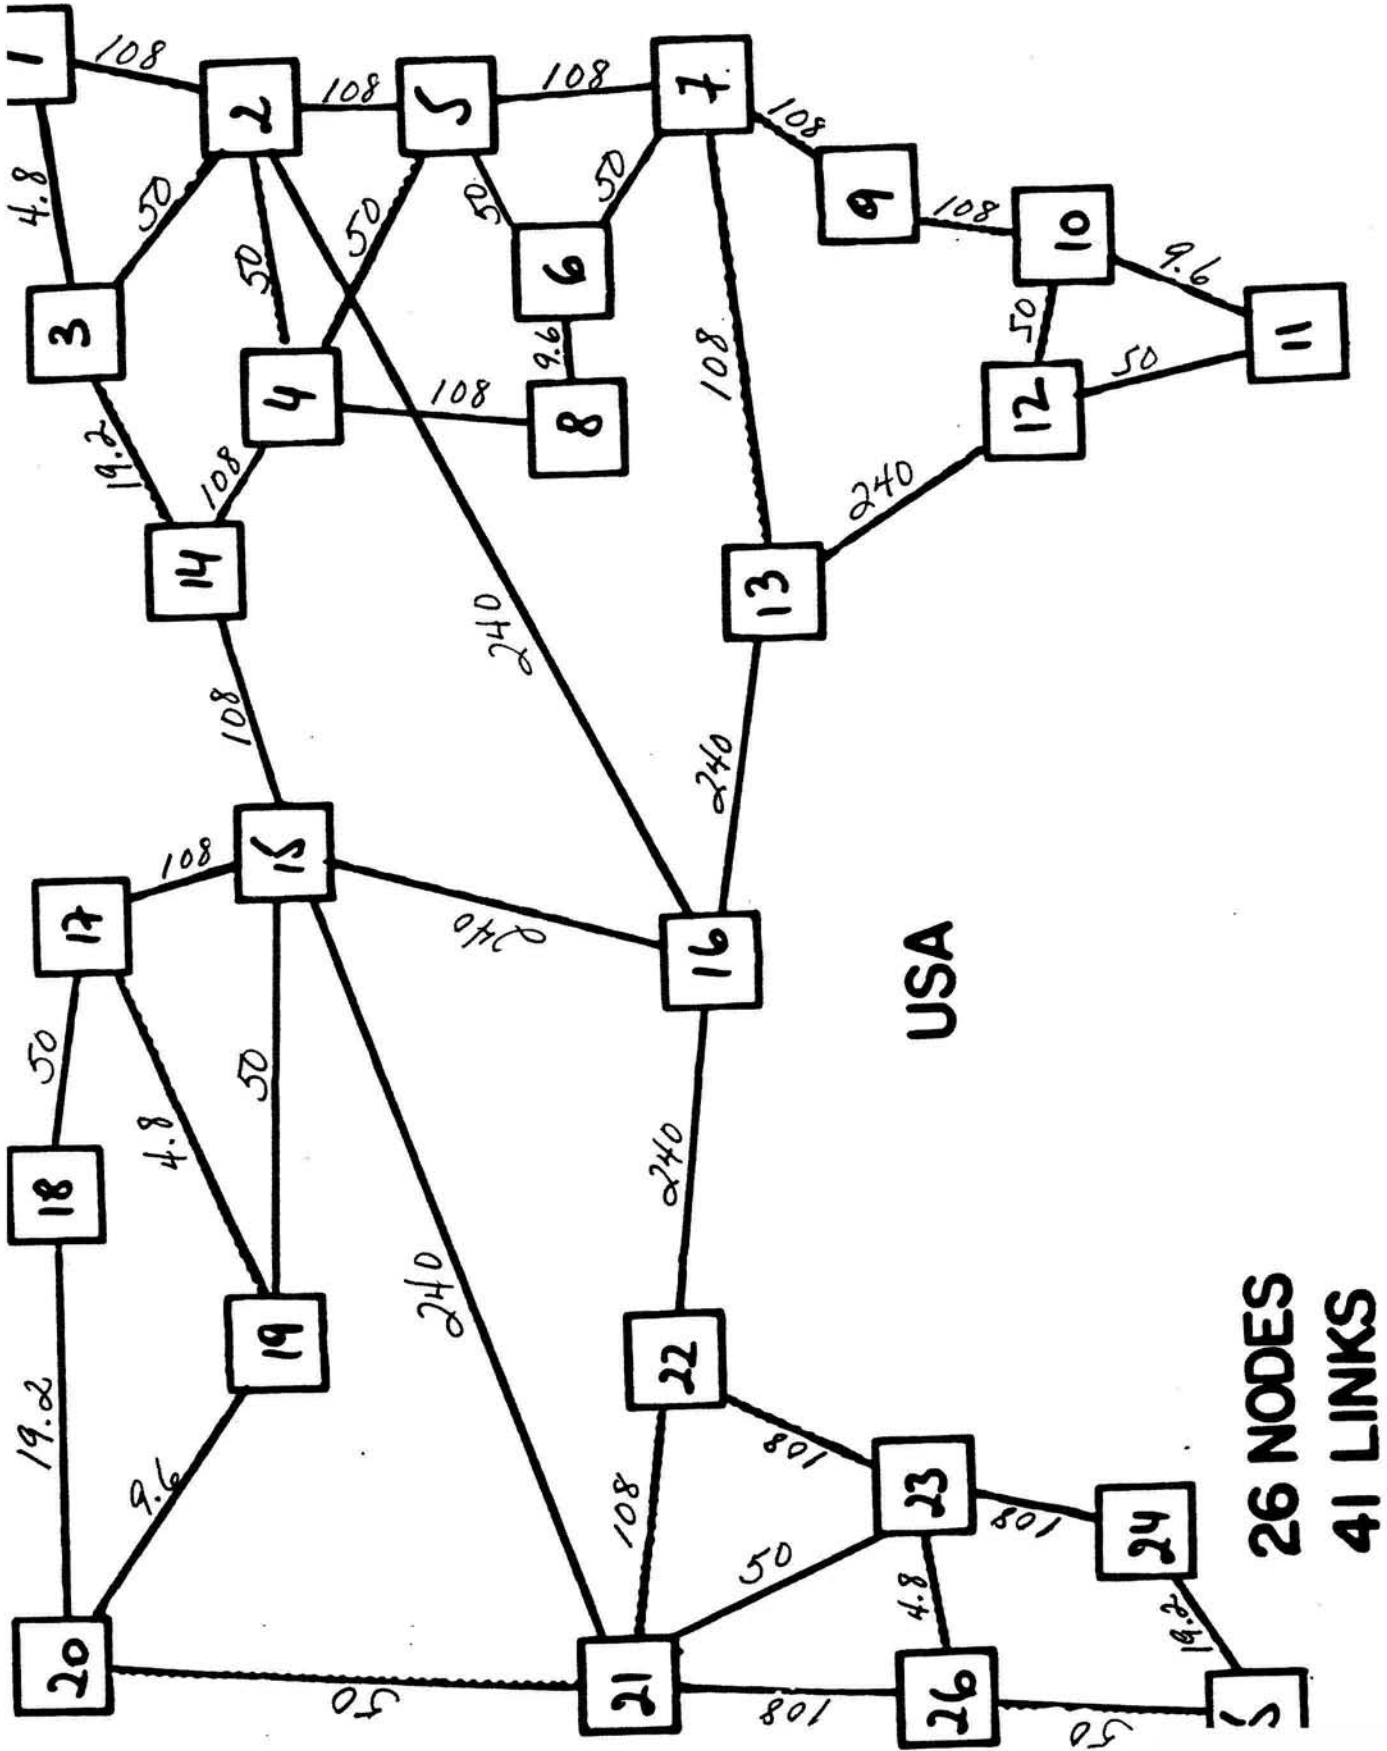



Figure 3: Topology and link capacities (Kbps) for the USA network

**26 NODES  
41 LINKS**



**RING**  
**32 NODES**  
**60 LINKS**

Figure 4: Topology and link capacities (Kbps) for the RING network

Network ID	Message length (cl. 1)	Lower bound	Upper bound	Cost cl. 1	Cost cl. 2	Upper/lower	Average delay (cl. 1)	Average delay (cl. 2)
GTE	400	19435	19447 (21501)	6780 (6429)	12667 (15072)	1.001	25.68 (24.35)	12.00 (14.27)
GTE	600	24648	24719 (24587)	12052 (11141)	12667 (13446)	1.003	45.65 (42.20)	12.00 (12.73)
GTE	800	32679	33417 (35486)	20705 (22105)	12712 (13381)	1.023	78.43 (83.73)	12.04 (12.67)
ARPA	400	10880	10882 (22368)	7815 (7045)	10267 (15323)	1.000	18.61 (16.77)	6.11 (9.12)
ARPA	600	25603	25619 (29007)	15351 (14170)	10268 (14837)	1.001	36.55 (33.74)	6.11 (8.83)
ARPA	700	32182	32434 (36352)	22153 (20577)	10281 (15775)	1.008	52.75 (48.99)	6.12 (9.39)
ARPA	800	44030	44914 (52232)	34630 (33547)	10284 (18684)	1.020	82.45 (79.87)	6.12 (11.12)
USA	400	43599	43908 (62851)	18712 (16149)	25196 (46702)	1.007	14.40 (12.42)	4.84 (8.98)
USA	600	62351	62837 (88013)	37632 (34292)	25205 (53721)	1.008	28.95 (26.38)	4.85 (10.33)
USA	800	108276	114177 (166574)	88709 (99112)	25468 (67462)	1.054	68.24 (76.24)	4.90 (12.97)
RING	400	100138	100180 (116422)	44234 (37631)	55946 (78791)	1.000	11.15 (9.48)	3.52 (4.96)
RING	600	155813	156309 (188609)	100157 (95176)	56152 (93432)	1.003	25.24 (24.00)	3.54 (5.89)
RING	700	217432	220312 (303858)	163336 (183184)	56976 (120674)	1.013	41.16 (46.16)	3.59 (7.60)
RING	800	382140	410223 (1392442)	352836 (283322)	57387 (1109120)	1.073	88.92 (71.40)	3.62 (69.88)

**Table 5:** Results for the priority and the FCFS case  
different message lengths ( $D_1=1000$ ,  $D_2=4000$ )

Network ID	Message lengths (cl.1/cl.2)	Lower bound	Upper bound	Cost cl.1	Cost cl.2	Upper/lower	Average delay (cl.1)	Average delay (cl.2)
GTE	400/400	40123	40183	10423	29760	1.001	39.48	28.18
GTE	600/200	24648	24719	12052	12667	1.003	45.65	12.00
GTE	700/100	19176	19583	13634	5949	1.021	51.65	5.63
ARPA	400/400	37357	37455	13460	23995	1.003	32.05	14.28
ARPA	600/200	25603	25619	15351	10268	1.001	36.55	6.11
ARPA	700/100	21862	22060	17250	4810	1.009	41.07	2.86
USA	400/400	96261	97212	34841	62371	1.010	26.80	11.99
USA	600/200	62351	62837	37632	25205	1.008	28.95	4.85
USA	700/100	52408	53419	41864	11555	1.019	38.20	2.22
RING	400/400	268693	269290	116683	152607	1.002	29.41	9.62
RING	600/200	155813	156309	100157	56152	1.003	25.24	3.54
RING	700/100	130685	131366	106450	24916	1.005	26.83	1.57

**Table 5:** Results for fixed total load ( $D_1=1000$ ,  $D_2=4000$ )

Network ID	Cost of delay (c1.1)	Lower bound	Upper bound	Cost c1.1	Cost c1.2	Upper/lower	Average delay (c1.1)	Average delay (c1.2)
GTE	4000	60607	60780	48078	12712	1.003	45.52	12.04
GTE	1000	24648	24719	12052	12667	1.003	45.65	12.00
GTE	100	13804	13946	1268	12678	1.010	48.04	12.00
GTE	1	12133	12715	34	12681	1.048	130.52	12.01
ARPA	4000	71365	71613	61164	10449	1.003	36.41	6.22
ARPA	1000	25603	25619	15351	10268	1.001	36.55	6.11
ARPA	100	11505	11811	1556	10255	1.027	37.05	6.10
ARPA	1	9736	10290	17	10273	1.057	41.37	6.11
USA	4000	174171	175215	149545	25670	1.006	28.76	4.94
USA	1000	62351	62837	37632	25205	1.008	28.95	4.85
USA	100	28557	29013	3817	25196	1.016	29.36	4.85
USA	1	23988	25235	39	25196	1.052	30.26	4.85
RING	4000	452967	455906	399395	56511	1.006	25.16	3.56
RING	1000	155813	156309	100157	56152	1.003	25.24	3.54
RING	100	65871	66002	10131	55871	1.002	25.53	3.52
RING	1	55018	56014	144	55870	1.018	36.40	3.52

**Table 5**: Results for different unit costs of delay ( $1/\mu_1=600, 1/\mu_2=200$ )

## References

- [Ahuja 79] V.Ahuja. Routing and flow control in systems networks architecture. *IBM Syst. Jour.* 18:298-314, 1979.
- [Amoss 80] J. J. Amoss. Planning for the Bell Operating Systems network. In *Proc. of the 5th ICC, Atlanta*, pages 559-563. 1980.
- [Dureste 83] B. Dureste and J. Kravitz. The SITA telecommunications network. In *IEEE Intern. Conf.*, pages 1067-1069. 1983.
- [Gavish 82] B. Gavish. Topological design of centralized computer networks: Formulation and algorithms. *Networks* 12:355-377, 1982.
- [Gavish 83] B. Gavish and S.L. Hantler. An algorithm for the optimal route selection in SNA networks. *IEEE Trans. Commun.* COM-31:1154-1161, 1983.
- [Gavish 86] B. Gavish and I. Neuman. Capacity and flow assignment in large computer networks. In *Proceedings, INFOCOM'86, Miami, Florida*, pages 275-284. April, 1986.
- [Held 71] M. Held and R.M. Karp. The travelling salesman problem and minimum spanning trees, Part II. *Math. Program.* 1:6-25, 1971.
- [Kleinrock 76] L. Kleinrock. *Queuing systems, vol. 2: Computer applications*. John Wiley & Sons, 1976.
- [Kroneberg 85] D. Kroneberg. A packet/message switch for the SITA network. In *ICC'75, San Francisco*. June 1985.
- [Laue 81] R.V. Laue. A versatile queuing model for data switching. In *Symp. on Data Commun.*, pages 118-128. 1981.
- [Maruyama 76] K. Maruyama and D.T. Tang. Discrete link capacity assignment in communication networks. In *3rd ICC, Toronto*, pages 92-97. 1976.
- [Maruyama 77] K. Maruyama, K. Fratta and D.T. Tang. Heuristic design algorithm for computer communication networks with different classes of customers. *IBM Jour. Res. Develop.* :360-369, 1977.
- [Morris 81] R.J.T. Morris. Priority queuing networks. *Bell Sys. Tech. Jour.* :1745-1769, 1981.
- [Neuman 86] I. Neuman. *Routing in a network with different classes of messages*. Technical Report, Center for Research on Information Systems, Computer Applications and Information Systems Area, Graduate School of Business Administration, New York University, 1986.



- [Pandya 77] R.N. Pandya. Delay analysis for DATAPAC: a packet switched network with two priority classes. In *5th Data Commun. Symp., Snowbird, Utah*, pages 314-121. 1977.
- [Reiser 79] M. Reiser. A queuing network analysis of computer communication networks with window flow control. *IEEE Trans. Commun.* COM-27:1199-1209, 1979.
- [Tang 76] K. Maruyama and D. T. Tang. *Discrete link capacity and priority assignment in communication networks*. Technical Report RC 6121, IBM Thomas J. Watson Research Center, Yorktown Heights, New York, 1976.