SOME RESULTS ON THE DESIGN OF COMMUNICATION NETS †

Leonard Kleinrock
Department of Engineering
University of California
Los Angeles, California 90024
Telephone: Area Code 213, 478-9711, Ext. 7202, 7212

Abstract

An approach to the design of communication nets is presented in which message delay is the performance index. An optimal channel capacity assignment is solved for in a capacity-constrained net, and its performance is discussed. The variation of message delay with topological structure and routing procedure is also investigated.

I. Introduction

The efficient design of a communication net is an extremely complex problem. The number of design parameters and operating modes is considerable and the definition of efficiency is by no means unique. Furthermore, the environment in which the net must operate strongly influences its construction. As a consequence, many discussions of communication network design tend to be either shallow qualitative treatments of rather general situations or detailed treatments of highly specialized and simplified nets. In this paper an attempt is made to give quantitative results for an interesting class of networks and then to discuss the implications of these results as certain of the other design parameters are varied. The philosophy of communication network design given here draws upon material presented in the author's book [1]. Other approaches have been studied in depth and many of these are summarized in the book by Ford and Fulkerson [2] and also in the forthcoming book by Frank and Frisch [3].

II. The Model

A communication net is made up of a collection of communication centers (nodes) which are connected together by a set of communication channels (ordered links). Messages (which are described by their origin, destination, origination time, length, and priority class) flow through the network in a store-and-forward fashion (as opposed to direct-wire or telephone traffic). A set of operating rules for handling the message traffic must also be given in describing a communication net.

† This research was supported in part by the United States Army Electronics Command, Contract DAAB07-67-C-0540; the Advanced Research Projects Agency, SD-184; and the Applied Mathematics Division, Office of Research, United States Air Force, Grant No. AFOSR-88-1408.

The primary function of the network is to provide rapid and reliable communication between many of its communication centers simultaneously. The design of such networks involves a number of operational aspects of the stochastic flow of message traffic: message routing procedures; priority queueing disciplines; channel capacity assignments; and topological configurations. Furthermore, the environment in which the net must operate may be highly variable and even hostile. In the following subsections we describe these operational procedures and give an indication of the form they should take in a variety of environments.

2.1 Message Routing Procedures

A message routing procedure is a decision rule which determines, according to some algorithm (possibly random), the next node which a message will visit. The specification of the algorithm specifies the routing procedure. The parameters involved in the algorithm may include such things as: origin and destination of the message; priority of the message; availability of certain channels; congestion (or annihilation) of certain nodes.

We define a fixed routing procedure as one in which a message's path through the net is uniquely determined once its origin and destination are given. If more than one path is allowed, then we consider this to be an alternate routing procedure. An alternate routing procedure may choose its alternate paths either deterministically or at random (from some appropriate distribution) from among the operating links based upon the parameter values mentioned above; the former may be referred to as deterministic alternate routing and the latter as random alternate routing (or more simply as random routing procedures). See [1, 4, 5].

2.2 Priority Queueing Disciplines

In passing through the net, messages are often required to form a queue while awaiting transmission between nodes, and very often a priority discipline describes the structure of the queue. Priority queueing refers to those disciplines in which an entering message is assigned a set of parameters (either at random or based upon some property of the message) which determine its relative position in the queue. This position may vary as a function of time owing to the appearance of higher priority messages in the queue. In attempting to meet the user's demands in regard to a priority structure, these demands, in their most useful form, must stipulate the following:

1. The number of priority classes.
2. The relative performance each class expects from the network (expressed in terms of average delay time).
3. The average number of messages arriving per second and the average message length for each priority class.

In addition, it would be helpful (but not necessary) if the cost as a function of message delay for the various priority classes were supplied by the user; this information would allow a careful determination of the necessary service given each class (item 2 above). It should be emphasized that even partial information here is useful; for example, if an extremely high relative cost is associated with the highest priority class, then some form of preemptive queue discipline [6] is almost surely called for. If the distribution of message arrival times and message lengths were available, this too would permit the development of a more realistic model for analysis.

With such a set of user demands, it is possible to design priority queueing disciplines with enough adjustable parameters so as to be able to comply with these demands (see, for example [7]).

In many communication systems, each channel may naturally be broken into a number of subchannels. The conditions under which this may be desirable must be carefully considered; such considerations include the cost (in terms of delay) of preemption, the particular priority class structure, etc. In a nonpriority case, it can be shown that one should never subdivide a channel if delay is the only criterion (see [1], Theorem 4.2).

2.3 Topology

The topological configuration of the communication net strongly affects its behavior as regards reliability, message delay, routing, etc. (see [1, 4, 5]). It is clear that complete freedom is not generally available in the design of the topology for most nets. In fact, not only may the topology be constrained, but also it may be that the structure of the network will be changing during the period of its operation. However, proper advantage must be taken of the freedom which does remain in the restructuring during configuration changes so that "optimal" performance is achieved.

Once the topological constraints have been met, there remains the crucial problem of selecting the capacity of each channel in the net, (the channel capacity assignment). We address ourselves to this problem in the following section.

III. Optimal Channel Capacity Assignment

We first consider a situation in which there are N separate single exponential channel facilities. The ith node has a Poisson arrival rate λ_i messages per second, each message having an exponentially distributed length of mean $1/\mu_i$ bits; the channel capacity associated with the ith node is C_i. All nodes behave independently of each other; however, they are mutually coupled by the following linear constraint on their capacities:

$$C = \sum_{i=1}^{N} C_i \tag{1}$$

That is, there is distributed throughout the N channels a total capacity of C bits/sec. The system under consideration is shown in Figure 1.

For any assignment of the C_i which satisfies Equation (1), there is defined† T_i = E (total time that a message spends in waiting for and passing through channel i). One may ask about that particular assignment of the C_i which satisfies Equation (1) and which also minimizes the average (over the index i) of the set of numbers T_i. Specifically, we define this average to be

$$T = \sum_{i=1}^{N} \frac{\lambda_i}{\lambda} T_i \tag{2}$$

where

$$\lambda = \sum_{i=1}^{N} \lambda_i \tag{3}$$

†E(x) denotes mathematical expectation of x.

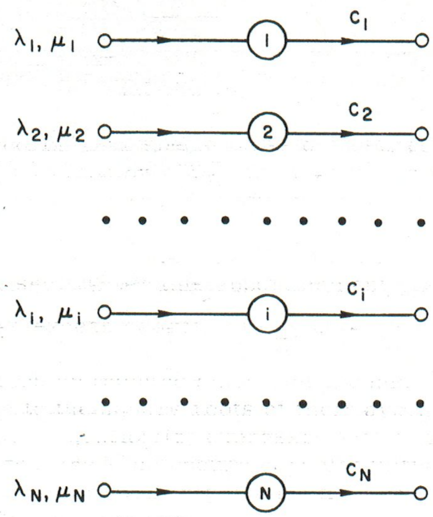

FIG I SYSTEM OF N SEPARATE SINGLE CHANNEL FACILITIES

Note that the weighting factor λ_i/λ for T_i has been chosen in the obvious way to be proportional to the number of messages which pass through node i. The solution to this problem is stated in:

Theorem 1

The assignment of the set C_i which minimizes T and which satisfies Equation (1) is

$$C_i = \frac{\lambda_i}{\mu_i} + C(1-\rho)\frac{\sqrt{\lambda_i/\mu_i}}{\sum_{j=1}^{N}\sqrt{\lambda_j/\mu_j}} \tag{4}$$

provided that

$$C > \sum_{i=1}^{N}\frac{\lambda_i}{\mu_i} \tag{5}$$

where

$$\rho = \frac{\lambda}{\mu C} \tag{6}$$

and

$$\frac{1}{\mu} = \sum_{i=1}^{N}\frac{\lambda_i}{\lambda}\frac{1}{\mu_i} \tag{7}$$

With this optimum assignment, we find that

$$T_i = \frac{\sum_{j=1}^{N}\sqrt{\lambda_j/\mu_j}}{C(1-\rho)\sqrt{\lambda_i\mu_i}} \tag{8}$$

and

$$T = \frac{\left(\sum_{i=1}^{N}\sqrt{\lambda_i/\mu_i}\right)^2}{\lambda C(1-\rho)} \tag{9}$$

This theorem is a special case of Theorem 2 which is proved below. We note that the optimum assignment operates in the following way. Each channel is first apportioned just enough capacity to satisfy its average required flow of λ_i/μ_i bits/sec. After this apportionment, there remains an excess capacity $C - \sum_{i=1}^{N}\lambda_i/\mu_i = C(1-\rho)$ which is then distributed among the channels in proportion to the square roots of their average flows λ_i/μ_i. Equation (5) expresses the obvious condition that there be enough capacity initially to satisfy the minimum requirements of the average flow in each node.

We now consider the more general case of an interconnected net with N channels subject to a fixed routing procedure. We assume the interarrival times and message lengths are independent random variables throughout the net (see the Independence Assumption [1]). Furthermore, the externally applied traffic is Poisson in nature. Consequently, we find that the interarrival times for message arrivals throughout the net are also Poisson (see [8]). This being the case, we intend to show that the optimum† channel capacity assignment for the net, with fixed total capacity C, is described by an equation similar to Equation (4). The interpretation of λ_i is, as before, the average arrival rate of messages to the ith channel; further, we take $\mu_i = \mu$ for all i. The average message delay T now must be carefully defined as

$$T = \sum_{j,k}\frac{\gamma_{jk}}{\gamma}Z_{jk} \tag{10}$$

where γ_{jk} = average number of messages entering network, per second, with origin j and destination k

$$\gamma = \sum_{j,k}\gamma_{jk}$$

Z_{jk} = average message delay for messages with origin j and destination k.

That is, T is appropriately defined as the overall average message delay, where the weighting factor for Z_{jk} is taken to be proportional to the number of messages which must suffer the delay Z_{jk}. For any pair jk, the quantity Z_{jk} is

†Optimum being interpreted as minimizing T, the average message delay.

composed of the sum of the average delays encountered in passing through each channel on the fixed route from node j to node k. If we break Z_{jk} into such components, and if we also form T by summing over the individual delays suffered at each channel in the net (instead of summing the delays for origin-destination pairs), we immediately see that

$$T = \sum_i \frac{\lambda_i}{\gamma} T_i \tag{11}$$

where clearly λ_i is the sum of all γ_{jk} for which the (fixed) jk route includes channel i. Thus we note that T is defined in a consistent manner [that is, $\lambda = \gamma$ for the net in Figure 1, and so Equations (2) and (11) are equivalent]. We may now state our fundamental result:

Theorem 2

For a net as described above, with a fixed routing procedure, the optimum channel capacity assignment is

$$C_i = \frac{\lambda_i}{\mu} + C(1-\bar{n}\rho)\frac{\sqrt{\lambda_i}}{\sum_{j=1}^{N}\sqrt{\lambda_j}} \tag{12}$$

With this optimum assignment,

$$T_i = \frac{\sum_{j=1}^{N}\sqrt{\lambda_i}}{\mu C(1-\bar{n}\rho)\sqrt{\lambda_i}} \tag{13}$$

and the average message delay T is

$$T = \frac{\bar{n}\left(\sum_{i=1}^{N}\sqrt{\lambda_i/\lambda}\right)^2}{\mu C(1-\bar{n}\rho)} \tag{14}$$

where $\bar{n} = \lambda/\gamma$ is the average path length for messages, and $\rho = \gamma/\mu C$.

Proof: From our assumptions of independence, each queue may be treated separately. From the well-known results from queueing theory (see, for example, [9]) we have for the ith channel with Poisson arrival rate λ_i and exponential service rate $\mu_i C_i$ that the delay T_i is given by

$$T_i = \frac{1}{\mu_i C_i - \lambda_i}$$

We now wish to minimize T in Equation (11) subject to the capacity constraint in Equation (1). To this end we form the Lagrangian [10]

$$G = T + \alpha\left(\sum_{i=1}^{N} C_i - C\right)$$

where α is some undetermined multiplier. Forming $\partial G/\partial C_i = 0$ we get

$$C_i = \frac{\lambda_i}{\mu_i} + \frac{1}{\sqrt{\alpha\gamma}}\sqrt{\frac{\lambda_i}{\mu_i}} \tag{15}$$

Summing this last on i and using Equation (1) we get

$$\sum_{i=1}^{N} C_i = C = \sum_{i=1}^{N}\frac{\lambda_i}{\mu_i} + \frac{1}{\sqrt{\alpha\gamma}}\sum_{i=1}^{N}\sqrt{\frac{\lambda_i}{\mu_i}}$$

Solving this for $\sqrt{\alpha\gamma}$ and substituting back into Equation (15) we obtain Equation (12) (for $\mu_i = \mu$). Using this value for C_i in the expression for T_i and T gives us Equations (13) and (14). We note that this establishes Theorem 1 (where $\bar{n} = 1$).

It remains to prove that the average path length $\bar{n} = \lambda/\gamma$. We observe that

$$\bar{n} = \sum_{j,k}\frac{\gamma_{jk}}{\gamma} n_{jk} \tag{16}$$

where n_{jk} is the path length for the origin-destination pair jk. Now, we recognize that λ_i is the sum of all γ_{jk} for which the jk route includes channel i. If we consider λ (the sum of the λ_i), we observe that it is composed of the sum of the numbers γ_{jk} each added in n_{jk} times. Thus

$$\lambda = \sum_i \lambda_i = \sum_{j,k}\gamma_{jk} n_{jk} \tag{17}$$

Thus, from Equations (16) and (17) we get

$$\bar{n} = \lambda/\gamma$$

We note that $\rho\left(= \frac{\gamma/\mu}{C}\right)$ is the ratio of the average number of bits/sec. entering the net (from external sources) to the total bit-handling capacity $C = \sum_{i=1}^{N} C_i$. We thus refer to ρ as the network load. On the other hand, λ is the average number of messages being transmitted per sec. internal to the net. Since $\bar{n} = \lambda/\gamma$, we see that $\bar{n}$ is thus the ratio of the average number of messages being transmitted/sec. internally to the average number/sec. entering the net.

We observe that the assignment in Equation (12) operates in the same way as in Equation (4). This is discussed below.

IV. Message Delay and Network Design

The single most significant performance measure of a communication net (operating in a relatively stable, peaceful environment) is the average time, T, that a message spends in the net. Indeed, the assignment given by Equation (12) is that which minimizes T, and the value it takes on is given by Equation (14). This equation reveals the trade-off between two crucial properties of the net, namely, the average path length, $\bar{n}$, and the degree to which the traffic is concentrated (see below). We consider that the network designer is given values of μ, C, and γ (and therefore, also the network load ρ). He has available, as design parameters, the routing procedure, the channel capacity assignment, the topology and the priority discipline. Equation (14) assumes that the C_i have been optimally chosen according to Equation (12) and that a given topology and fixed routing procedure are in effect. However, since any fixed routing procedure and any topology applies, the designer may choose these in a way which minimizes T.

The numerator sum in Equation (14) is a convex function of λ_i/λ and thus attains its minimum value when $\lambda_i = \lambda$ for some $i = i_0$ and $\lambda_i = 0$ for $i \neq i_0$ subject to the condition $\sum_{i=1}^{N} \lambda_i/\lambda = 1$. Since the λ_i depend upon the input traffic (not a design variable) and the topology and routing procedure (clearly design variables), we should attempt to concentrate the λ_i as much as possible in order to minimize the numerator sum in Equation (14). Complete concentration is not possible since each node may be required to serve as both an origin and destination for some message traffic; subject to this, however, the net which achieves a maximal concentration of traffic is the star net, shown in Figure 2. We see

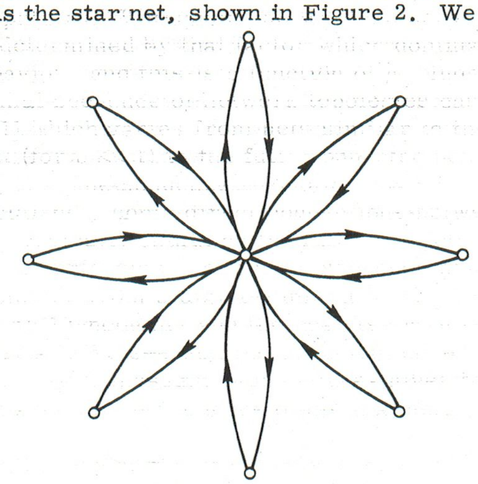

FIG 2 THE STAR NET

that this configuration groups all traffic leaving a node into a single channel (and likewise for the traffic entering a node) with the exception of the central node. We note that $\bar{n} \approx 2$ for the star net.

On the other hand, the denominator in Equation (14) contains the term $(1-\bar{n}\rho)$. As $\rho \to 1/\bar{n}$ we see that this term dominates the behavior of T, and so we must take care to minimize $\bar{n}$ in such a circumstance. The net which minimizes $\bar{n}$ is the fully connected net shown in Figure 3. For this net, all paths are of length unity and so $\bar{n} = 1$ (its minimum value).

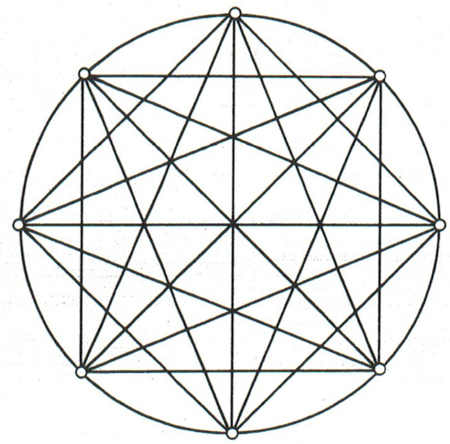

FIG 3 THE FULLY CONNECTED NET

It is clear that these two factors (i.e., the numerator sum which reflects the degree of traffic concentration, and the average path length $\bar{n}$) cannot be minimized independently. The trade-off between these two is apparent: the maximally concentrated traffic pattern (star net) yields an $\bar{n} \approx 2$; the (fully connected) net which minimizes $\bar{n}$ results in a maximally dispersed traffic pattern. The choice as to which net to use is determined by that factor which dominates the behavior, and this is a function of ρ. Indeed, an optimal sequence of network topologies can be found [1] which varies from nets similar to the star net (for $\rho \ll 1$) to the fully connected net (as $\rho \to 1$). The sequence is obtained by adding, to the star net, some direct connections between nodes, eventually obtaining the fully connected net. The performance of such a sequence has been obtained from a digital simulation [1] for a particular 13-node net and the results are given in Figure 4. We see that the minimum envelope of the family of message delay curves gives the performance of the sequence of optimal nets.

Having discussed the effect of topology and channel capacity assignment on message delay, we now direct our attention to the effect of

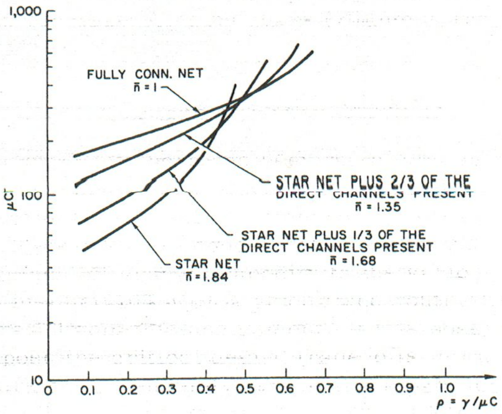

FIG 4 MESSAGE DELAY FOR THE SEQUENCE OF OPTIMAL NETS AS OBTAINED BY SIMULATION OF A 13-NODE NET

alternate routing procedures. Equation (12) assumes that the message traffic is stationary in time and also that a fixed routing procedure is used. It is clear that the communications requirement for many communication networks may change during their period of operation (especially if the environment is fluid and/or hostile). Consequently, the network must be capable of dynamically altering its characteristics during such changes. In that case, one cannot hope to maintain a fixed channel capacity assignment such as that given in Equation (12) and a fixed routing procedure. Indeed, if the input traffic requirements are changing, then the numbers λ_i cannot be calculated and C_i cannot be chosen according to Equation (12). In such a circumstance, a judicious alternate routing procedure serves a major role. Indeed, such a routing procedure automatically matches the traffic pattern to the network topology. Furthermore, it tends to distribute the traffic so that the channel capacity assignment which happens to exist at any time appears as one in which the traffic and capacity are proportional (see [1]). This resultant proportional channel capacity assignment is clearly not the optimum (indeed, that given by Equation (12) is the optimum), however, it is nearly optimum, and is therefore desirable. Moreover, as the environment becomes more hostile and/or mobile, directory-type routing procedures require more and more of the communication capability to update the rapidly changing directory (see [4, 5, 11]); in such cases, it is important to rely less upon exact and complete directories and to depend more upon local network information making use of alternate routing. A degree of caution must be exercised, however, since uncontrolled alternate routing in a congested net can lead to chaos. Indeed, the telephone company tends to limit (and even prohibit completely) alternate routing on unusually busy days (Mother's Day, for example).

V. Conclusions and Extensions

The emphasis here has been the minimization of the average message delay in a capacity-constrained communication network. The approach taken was to concentrate upon the optimal assignment of channel capacity to the various links in a network with a given fixed routing procedure. From this assignment, it was possible to expose the critical design trade-offs in the network design, namely, the average path length and the degree of concentration of the traffic. From these observations, it was then possible to derive an optimal sequence of topologies as a function of the network load. Indeed, as shown in [1], in a communications environment for which the terminal traffic conditions are known and time-invariant, the most efficient network design incorporates a topology chosen from the optimal sequence (dependent upon ρ), uses the channel capacity assignment given by Equation (12), and follows a fixed routing procedure. However, if the required traffic between nodes is either unknown or time-varying, then some form of alternate routing is essential. This statement follows from the observation that alternate routing procedures dynamically assign the traffic in a way which matches the message flow to the current network capabilities.

In the previous discussion, a fixed cost constraint (equal to the total channel capacity C) was assumed. Other cost functions which depend upon the channel capacity in a linear fashion have also been handled [1]. Cost functions which are nonlinear with the capacity suggest the use of iterative solution methods (e.g., steepest descent) for optimization. It is clear that many other factors enter into the performance of a communications net, for example, the network vulnerability to attack or failure. Considerations such as these require a new problem formulation; however, it appears that the results reported here give an indication of the general form and operation of networks optimized for a variety of situations.

References

1. L. Kleinrock, Communication Nets; Stochastic Message Flow and Delay, McGraw-Hill (1964).

2. L.R. Ford Jr. and D.R. Fulkerson, Flows in Networks, Princeton University Press, Princeton, N.J. (1962).

3. H. Frank and I.T. Frisch, Communication and Transmission Networks, to be published.

4. R.T. Prosser, "Routing Procedures in Communication Networks, Part I: Ransom Procedures," IRE Trans. on Communication Systems, CS-10 (4), pp. 322-329 (1962).

5. R.T. Prosser, "Routing Procedures in Communication Networks, Part II: Directory Procedures," IRE Trans. on Communication Systems, CS-10 (4), pp. 329-335 (1962).

6. H. White and L.S. Christie, "Queueing with Preemptive Priorities or with Breakdown," Operations Research, Vol. 6, pp. 79-95 (1958).

7. L. Kleinrock and R.P. Finkelstein, "Time Dependent Priority Queues," Operations Research, Vol. 15, No. 1, pp. 104-116 (1967).

8. P.J. Burke, "The Output of a Queueing System," Operations Research, Vol. 4, pp. 699-704 (1956).

9. T.L. Saaty, Elements of Queueing Theory with Applications, McGraw-Hill, New York (1961).

10. F.B. Hildebrand, Methods of Applied Mathematics, Prentice-Hall, Inc., Englewood Cliffs, N.J. (1958).

11. P. Baran, "On Distributed Communication Networks," IEEE Trans. on Communications Systems, pp. 1-9, March 1964.