Communications Research Centre

THE PLANNING OF DOMESTIC
COMMUNICATION NETWORKS

by

G.A. Neufeld, R.R. Bowen and A.R. Kaye

DEPARTMENT OF COMMUNICATIONS
MINISTÈRE DES COMMUNICATIONS

CRC TECHNICAL NOTE NO. 663

CANADA

OTTAWA, JANUARY 1974

COMMUNICATIONS RESEARCH CENTRE

DEPARTMENT OF COMMUNICATIONS
CANADA

THE PLANNING OF DOMESTIC COMMUNICATION NETWORKS

by

G.A. Neufeld, R.R. Bowen and A.R. Kaye

(Communications Systems Directorate)

CRC TECHNICAL NOTE NO. 663

January 1974
OTTAWA

TABLE OF CONTENTS

THE PLANNING OF DOMESTIC COMMUNICATION NETWORKS*

by

G.A. Neufeld, R.R. Bowen and A.R. Kaye

ABSTRACT

Techniques for the synthesis of domestic long-haul communication systems to meet forecast traffic requirements must provide for cost-effective integration of satellites with existing and new terrestrial systems. A methodology that may be used to carry out this synthesis is described. One of its important features is the separate investigation of the characteristics of satellite and terrestrial systems with merging of the results towards the end of the study. This allows activities to be carried out concurrently rather than consecutively so that both sets of results may be more up-to-date and so more useful when they are merged. Optimization algorithms are described which are of general interest in that they can be used to optimize any network having step-like link costs.

1. INTRODUCTION

This paper discusses a methodology which has been developed for the identification of cost-effective plans for the development of the Canadian long-distance communication networks for the 1980's. In particular, attention is focussed on the set of options and trade-offs to be investigated in planning the integration of satellite and terrestrial systems to meet forecast traffic requirements.

The nature of any satellite system which is introduced is a function of:

(a) the amount of trunk traffic between the various regions of Canada in the 1980's;

(b) the amount of terrestrial trunk communications equipment that is installed prior to 1980, or which can be installed as an

*The information contained in this note was presented in part at the 44th National ORSA Conference, San Diego, California, November 1973.

'add-on' to an existing system at small cost;

(c) the cost of new terrestrial and satellite systems in the 1980's.

A large number of parameters must be considered in selecting a satellite system. A major difficulty is the determination of the value of each parameter to yield the most cost-effective system. To resolve this difficulty, it is necessary to have a methodology which permits the analyst to determine the best value for each parameter in a systematic manner.

The object of this paper is twofold. The first objective is to show how to choose an optimum satellite system by using:

(a) a set of curves that specify the amount of traffic that would be sent through a satellite, as a function of the cost of sending it through the satellite;

(b) a set of curves that show the costs of providing a satellite system, as a function of the amount of traffic to be carried by it.

The set of curves in (a) will be called the demand curves and the set of curves in (b) will be called the cost curves. The combination of (a) and (b) will henceforth be referred to as the cost-demand strategy.

The second objective is to describe the network model and the optimization algorithms that have been developed in order to generate the above demand curves. The generation of the cost curves is an exercise in satellite system engineering and is not discussed in detail.

2. THE COST-DEMAND STRATEGY

In this section we describe the cost-demand strategy. We will describe, in turn, the demand and cost curves. Then we will show how the two sets of curves may be used in order to determine the optimum integrated system.

2.1 DEMAND CURVES FOR SATELLITE SYSTEM USE

Suppose that Canada has been divided into a number of 'traffic centres' between which it is necessary to provide long-haul heavy-route communications paths and that the assumptions below are valid.

(a) Forecasts are available for the traffic flow between all the traffic centres. All traffic flow is point-to-point in that each traffic requirement can be considered to have a single source and a single destination. The basic unit of traffic flow will be referred to as a voice-circuit.

(b) The capacity and annual operating cost are known for terrestrial transmission systems which exist at the beginning of the period for which the planning is being done and the cost of augmenting these systems is known.

(c) The cost of installing, operating and maintaining new systems during the period for which the planning is being done is known.

If we assume some particular cost of routing a single voice-circuit through the satellite system we have enough information to define a multi-commodity network over which the traffic must be routed at minimum cost. In the next section we will characterize this network problem and describe an algorithm to find minimal total cost solutions. Although this algorithm will specify the optimum routing strategy, what is wanted at this point is simply the amount of traffic to be carried by the satellite.

The satellite system demand curves are a description of what traffic would be sent through the satellite system as a function of the cost of sending it through the satellite. The satellite, its earth stations, backhaul, etc., are represented simply by an annual cost per voice-circuit through the satellite. The algorithm can be repeated for different values of annual cost through the satellite. The resulting relation between total traffic carried by the satellite and annual cost per satellite circuit is the satellite system demand curve. One possible curve is shown in Figure (1).

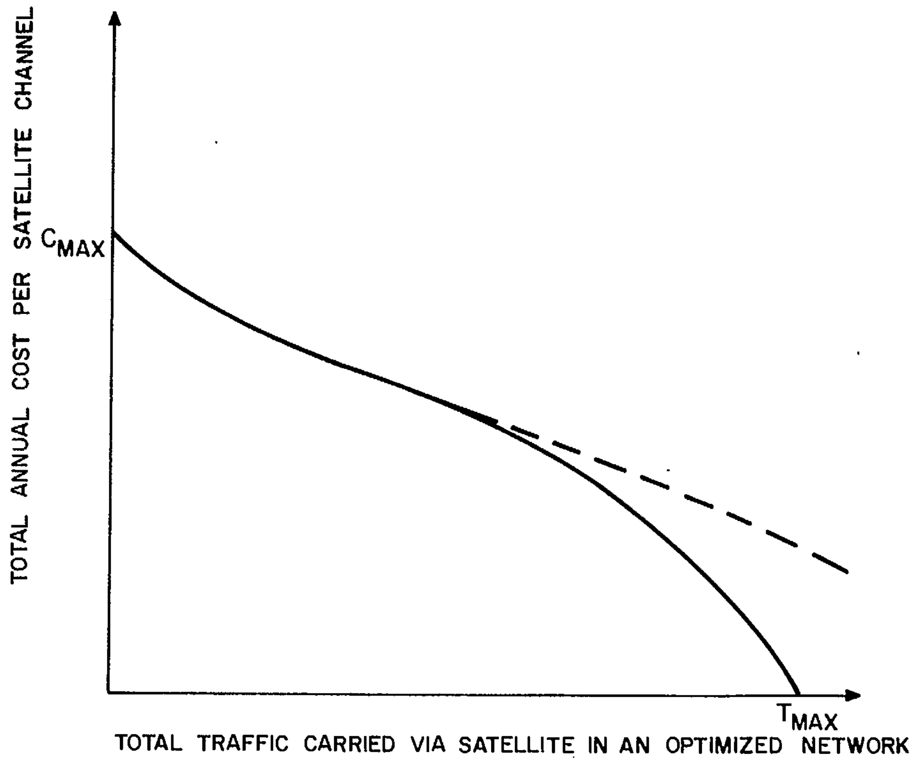

Fig. 1. Satellite system demand curve.

No account is taken in this analysis of the elasticity in total demand for trunk traffic as a function of cost. The total demand is assumed to be fixed and the solid line in Figure (1) simply shows the trade-off between terrestrial and satellite traffic as the cost of a satellite circuit is varied. If elasticity were to be accounted for, the demand curve would likely be of the form shown as a dotted line in Figure (1).

One property of Figure (1) is that the traffic through a satellite in an optimized system increases monotonically as the cost of a satellite channel is decreased. A second property is that at some cost C_{MAX} the traffic is zero, because at this cost it is more economical to use the terrestrial system exclusively. A third property is that the traffic cannot exceed T_{MAX} even though the satellite cost is zero, for reasons explained in the previous paragraph.

Suppose that we now consider a satellite system constrained to have not more than two earth stations, with no constraints on the position of those stations. (The necessity for this approach will become evident when the cost curves are discussed in Section 2.2.) A demand curve can be determined for this constrained system in much the same way that it was for the unconstrained one (Figure (2)). At very low levels of traffic through the satellite, or high satellite cost levels, it is likely that the demand curve of the constrained

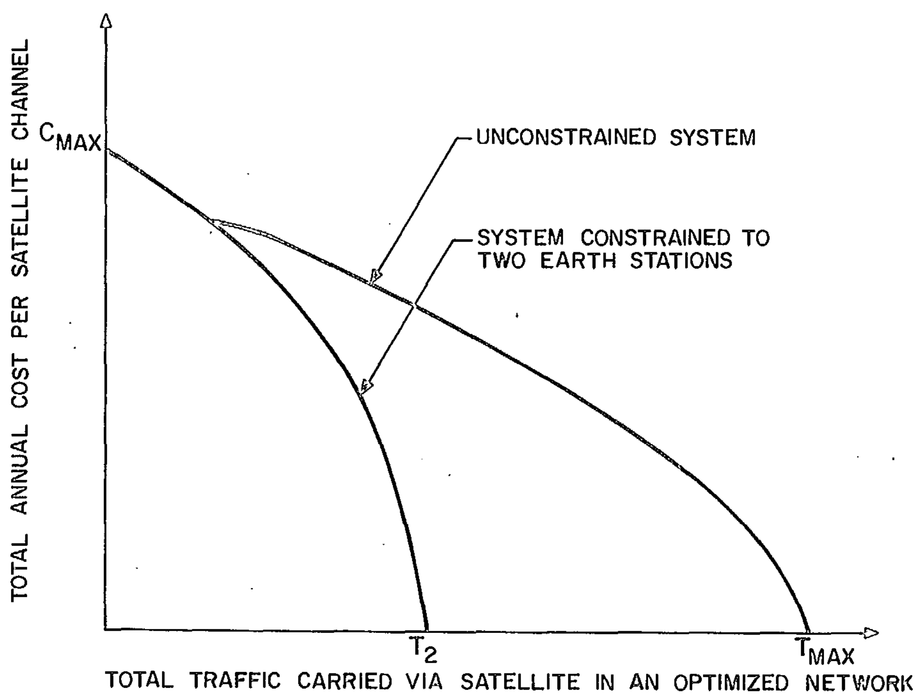

Fig. 2. Satellite system demand curves for the unconstrained system and a constrained one.

and unconstrained systems will be equal, since the unconstrained system would likely have only two earth stations. At higher traffic levels the constrained system will carry less traffic than the unconstrained one, because the unconstrained system would introduce more than two earth stations as the cost was decreased. Furthermore the demand curve of the constrained system would have a maximum traffic load T_2 below T_{MAX} because it would not be cost-effective to route T_{MAX} to the two earth stations. Traffic between other regional centres will undoubtedly be re-routed to take advantage of the low satellite link cost, but it is very unlikely that the satellite would take all such traffic. Typical curves for the constrained and unconstrained systems are shown in Figure (2).

The same argument can be made for a satellite system constrained to have three earth stations, four earth stations, etc. The result is a set of traffic demand curves, as shown in Figure (3).

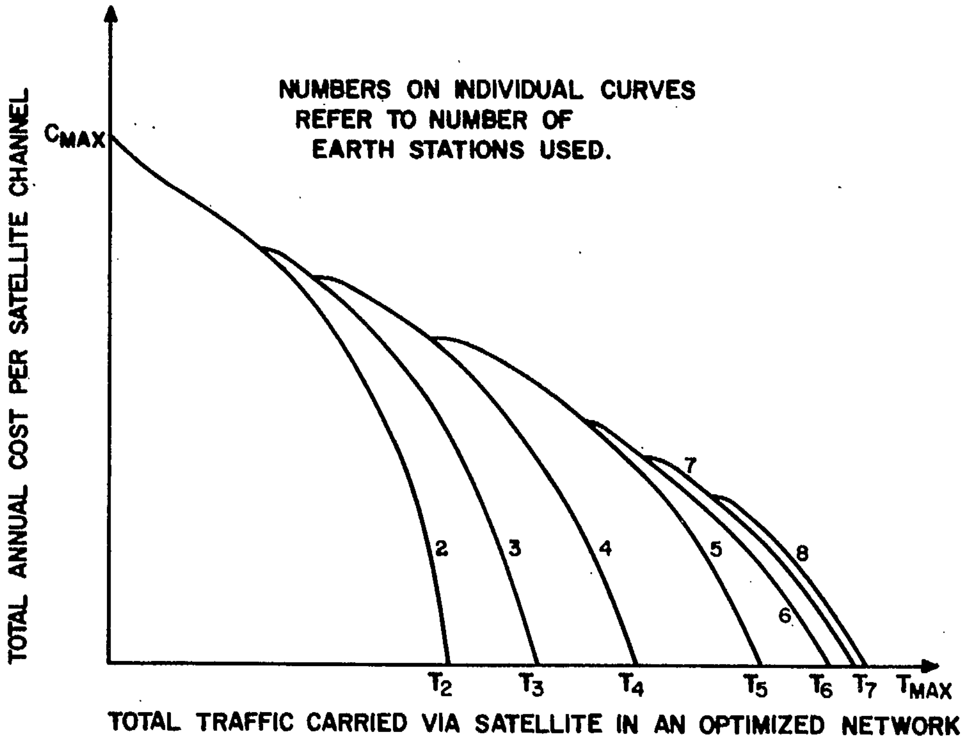

Fig. 3. Satellite system demand curves for various constrained systems.

The demand for the satellite in the unconstrained system will be greater than that for the constrained systems. However, on this unconstrained system curve we can separate the regions where that system has two, three, four, etc., earth stations. The unconstrained system curve will be the same as the constrained curve for n earth stations in the region in which the unconstrained optimum system uses n earth stations.

It is reasonable to suppose that the number of stations increases monotonically as the cost per satellite channel decreases, i.e., as the traffic through the satellite is increased. Consequently the demand curve for the unconstrained system is the envelope or upper bound of the curves for the constrained system. In generating a point on a demand curve the network algorithm specifies not only the total traffic carried by the satellite but also the location of each earth station and the traffic routed through each station. Thus each point on the curve is associated with a complete network synthesis.

2.2 SATELLITE SYSTEM COST CURVES

In generating the demand curves it was not necessary to know anything about the satellite system except the number of ground stations. In order to generate the cost curves, however, it is necessary to determine the optimum (minimum cost) design of a satellite system for each number of earth stations at all possible traffic levels. In doing this it is necessary to take into account all costs involved including maintenance and operation. Cost per circuit may be expressed in terms of present value of total charge or annual cost, whichever was the basis used in specifying the costs of the terrestrial network on which the demand curves were based. This procedure involves thorough engineering studies which are not the subject of this report. The result is a set of satellite system cost curves which might look like the hypothetical set shown in Figure (4). The cost curve for n earth stations, n = 1,2,3.., need not be determined for traffic levels higher than the maximum traffic level T_n specified by the demand curve for n earth stations.

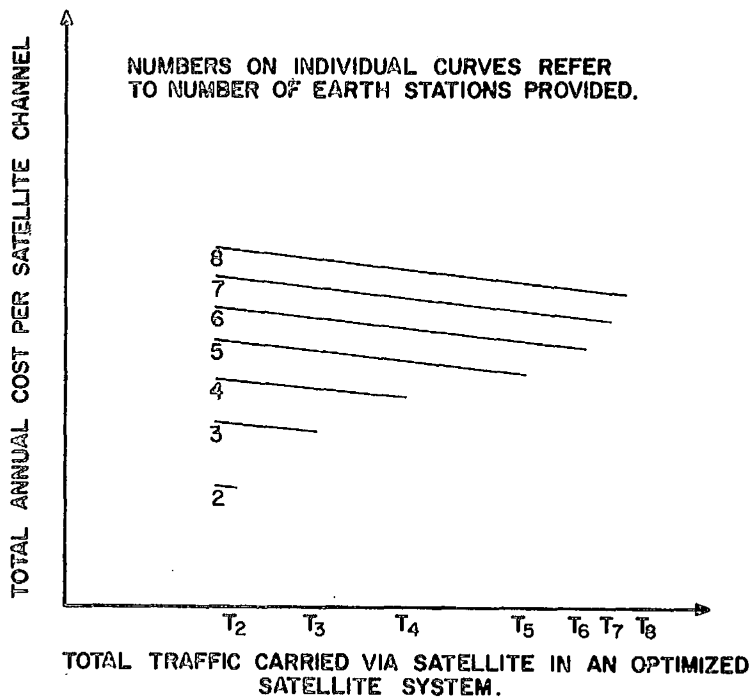

Fig. 4. Satellite system annual cost curves.

2.3 CHOICE OF THE GLOBAL OPTIMUM SATELLITE SYSTEM AND TRUNK NETWORK

So far we have obtained:

(a) a large number of 'optimum' networks, one for each assumed annual satellite circuit cost and number of earth stations;

(b) a large number of 'optimum' satellite system designs, one for each traffic distribution and number of earth stations.

Suppose we assume that the system has two satellite earth stations. We can plot the two-earth-station satellite-system demand curve and the two-earth-station satellite-system cost curve on the same graph. The point where the two curves cross corresponds to the optimum network and the optimum satellite system if only two earth stations are allowed, as shown in Figure (5).

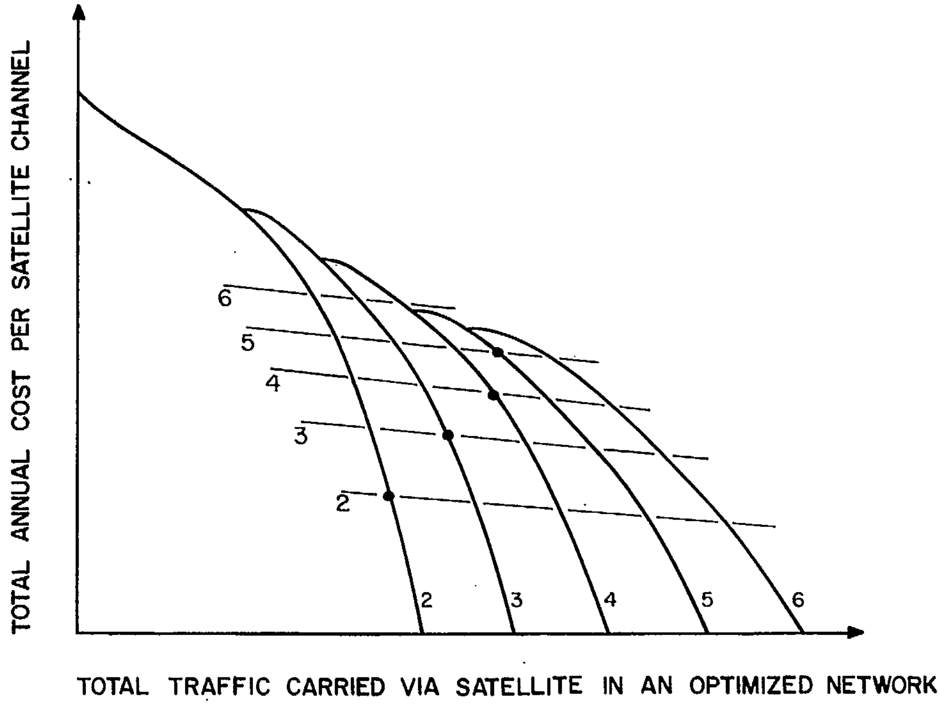

Fig. 5. Graphical determination of the global optimum system.

The same procedure can be repeated for a three-earth-station constraint, a four-earth-station constraint etc. Suppose that n earth stations is the largest number for which the satellite system cost curve and demand curve intersect, i.e., for (n+1) earth stations the cost curve is above the demand curve for all traffic values. Then the overall optimum system should have n earth stations, and the intersection of the two n-earth-station curves indicates the optimum total network. Even if there is a smaller number of earth-stations which would yield a lower average cost per satellite circuit the network with n earth-stations yields a lower overall network cost. In the example shown in Figure (5) the system should use five earth stations.

3. NETWORK MODEL AND ALGORITHM

In this section we first describe the network model that is required to generate the demand curves described above. From this, we define the network problem for which minimum cost solutions must be obtained. Then the algorithm that has been developed for finding solutions to the network problem is described.

3.1 NETWORK MODEL

The Canadian terrestrial network for long-distance communications consists of several transmission systems which are operated by the common carriers. Each of these transmission systems can be broken down into subsystems interconnecting two end points where an end point is either a source-sink for traffic or a 'junction point' where several subsystems meet. We have modelled the terrestrial communications network as a graph in which the nodes represent either a source-sink for traffic or a junction point. A link in the graph may represent either an existing or a possible new subsystem. There may be multiple links joining pairs of nodes. The graph of an interconnected terrestrial and satellite system must include an additional node that represents the satellite system. Links, which are incident to this 'satellite' node and any other node, represent paths for traffic flowing through the satellite. Figure (6) shows a simplified model of this network.

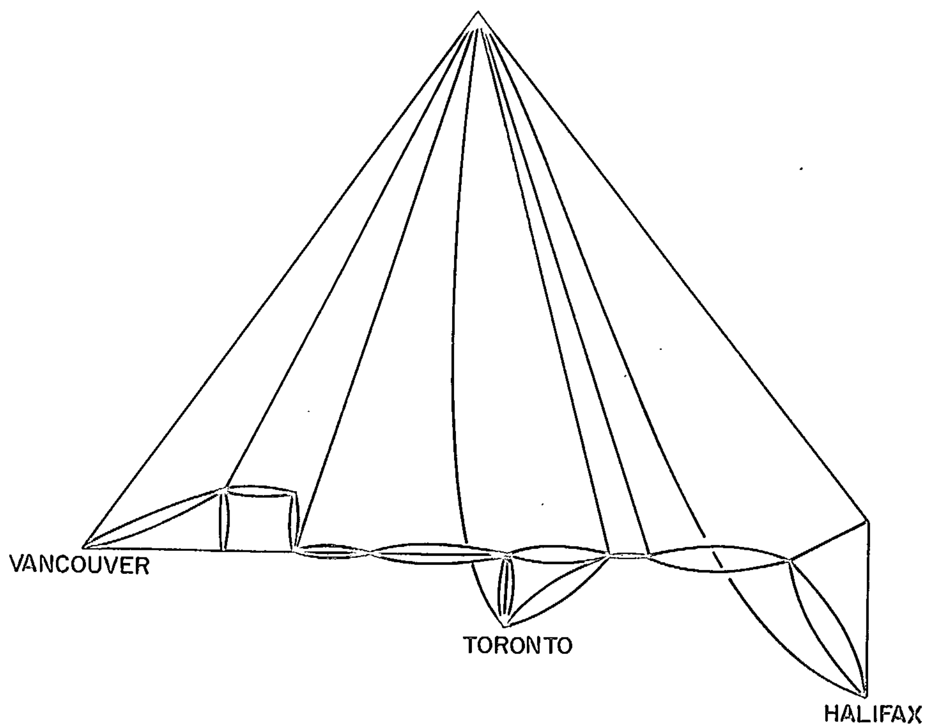

Fig. 6. Canadian network model.

Consider the links that represent the terrestrial transmission subsystems. Each one has a cost that is a function of the flow routed through, and hence the capacity required in, that link. A typical link cost function is shown in Figure (7). The link costs are discontinuous and piece-wise linear. The 'steps' of the cost functions represent the cost of initially installing or augmenting a subsystem. In the actual terrestrial system, traffic-dependent costs are incurred in the nodes themselves. On the other hand, routing algorithms become much more difficult to implement in networks which include both link and node costs. To avoid this problem the node costs are distributed among the links incident to each node.

Any traffic flowing through the satellite system must flow through two links connected to the 'satellite' node. Thus the cost C of routing a unit of traffic through the satellite system may be represented by assigning a cost of C/2 for each unit of traffic flowing through each of two links connected to the 'satellite' node. Thus we maintain a network with only link costs rather than both link and node costs.

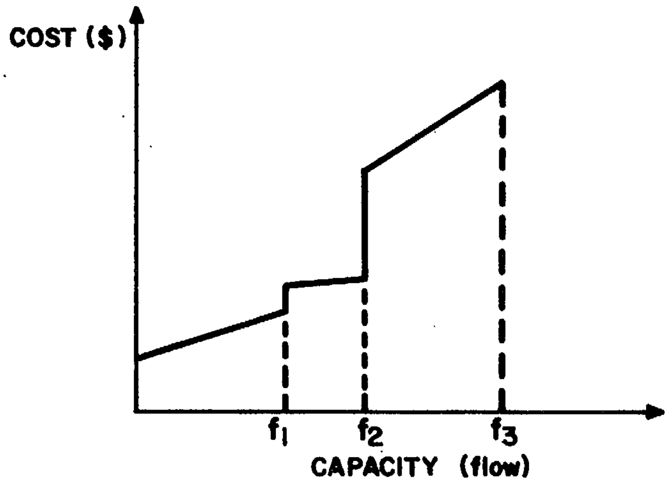

Fig. 7. Link cost function.

In summary, we have a multicommodity network flow problem and the objective is to route the traffic requirements through the network at minimal cost. The problem differs from the classical multicommodity flow problem in that the link costs are discontinuous and piece-wise linear rather than strictly linear.

More specifically, we consider an undirected graph G with nodes 1, 2, ...N and links $\ell_1,....\ell_M$. We can assume that there is only one link joining a given pair of nodes; intermediate nodes and links with zero cost can be introduced so as to eliminate multiple links between a pair of nodes. There are K commodities, where each commodity k has a source node s_k and a sink node t_k. Each commodity k has a flow requirement equal to f_k. Jewell's approach[4] to formulating problems with undirected links is applicable. Hence, let X^+_{km} be the flow of commodity k in link ℓ_m in one direction. Let X^-_{km} be the flow of commodity k in link ℓ_m in the opposite direction. Each link ℓ_m has capacity C_m where $C_m \geq 0$. It is assumed that there are enough links with sufficiently large capacities that a feasible solution exists. Let X_m be the total net flow through link ℓ_m. Each link ℓ_m has a step-like cost function $F_m(X_m)$ as in Figure (7). The steps need not have equal slopes nor equal lengths. The problem to be solved is

$$\text{minimize } Z = \Sigma_m F_m (X_m) \quad , \qquad .....(1)$$

subject to

$$X_{km} = X_{km}^{+} + X_{km}^{-} \quad \text{(all k and m)} \ , \qquad .....(2)$$

$$X_{km}^{+} \geq 0, \ X_{km}^{-} \geq 0 \ \text{(all k and m)} \ , \qquad .....(3)$$

$$X_m = \Sigma_k \ |X_{km}| \quad \text{(all m)} \ , \qquad .....(4)$$

$$X_m \leq C_m \quad \text{(all m), and} \qquad .....(5)$$

$$\Sigma_{m\varepsilon A_i^+} X_{km}^{+} + \Sigma_{m\varepsilon A_i^-} X_{km}^{-} - \Sigma_{m\varepsilon B_i^+} X_{km}^{+} - \Sigma_{m\varepsilon B_i^-} X_{km}^{-} = \begin{cases} f_k \text{ if } i = s_k \\ 0 \text{ if } i \neq s_k, t_k \\ -f_k \text{ if } i = t_k \\ \text{(for all i).} \end{cases} \qquad .....(6)$$

A_i^+ is the set of links that are incident to node i and such that X_{km}^+ represents the flow of commodity k towards node i. A_i^- is similar to A_i^+ with X_{km}^+ replaced by X_{km}^-. B_i^+ and B_i^- correspond to A_i^+ and A_i^- respectively except that the links correspond to flow away from node i. This problem corresponds to the classical telpak problem when the graph G is complete.

3.2 A HEURISTIC ALGORITHM

In this subsection we describe a heuristic algorithm for finding solutions to the network problem stated above. We first relate our work to that reported in the literature.

The telpak problem is a multicommodity flow problem in a complete graph with link cost functions that are a certain step function multiplied by Euclidean distance in a plane. Rothfarb and Goldstein[8] give a mathematical programming treatment of the telpak problem for the case where all the flow requirements are to a single point. However, as pointed out by White[10], the success of mathematical programming approaches is generally limited. This is due to the dimensionality of the problems and the nonlinearities of the cost functions. Kleitman and Claus[5] and Frank and Chou[2] discuss heuristic methods for finding solutions to the general telpak problem. Their methods for the telpak problem have several limitations when applying them to the problem being considered in this report. First, the methods perform local transformations on the network, in order to reduce the cost, by adding and deleting links on the basis of the relationship between the link costs and Euclidean distance. This transformation is not effective in the absence of this relationship. Second, it may not be possible to add longlines or 'direct lines' between a node pair if the graph is not assumed to be complete. Furthermore, as the variation of the 'steps' (or the size of the discontinuities) and of the length the piece-wise linear portions of the cost functions increases, it becomes increasingly more difficult to consider routing a particular flow requirement elsewhere in order to reduce the cost. Rather, a more global approach is required that considers the interdependence of the routing performed in order to take advantage of any available economies of scale. On the other hand, the discontinuous link costs present difficulties in applying the general theory of non-linear optimization[3,12].

The algorithm given here is iterative. At each iteration, the algorithm routes a subset of the flow requirements through the network and stops when all the flow requirements have been routed. The strategy is more comparable to that used in heuristic graph coloring algorithms[1,7,9,11] than that used in most sub-optimal network routing strategies; it is comparable to graph coloring algorithms in that nodes (traffic requirements) are assigned (routed) to colors (routes) in an iterative manner until all the nodes (traffic) requirements have been assigned (routed).

At each iteration the algorithm must determine the subset of the flow requirements to be routed and where to route it. We consider first the problem of where to route any subset of the flow requirement not yet routed. There is a lower bound y_m on the cost per unit of flow through each link ℓ_m and this is used as an estimated cost for that link. The lower bound y_m depends upon the link cost function and the total amount of flow b_m routed through link ℓ_m during previous iterations; $b_m = 0$ for all links ℓ_m during the first iteration. Ideally, each flow requirement would be routed along the 'shortest path' between the source and sink nodes in the network with each link ℓ_m having a weight equal to y_m. The weights y_m are however only estimated costs. The real cost per unit of flow through each link is dependent upon the amount of flow routed through the link. To illustrate, we consider a link with a cost function as in Figure (7) and with $b = 0$. This link would need to carry f_2 units of flow in order for the cost, per unit, to be equal to the least possible cost. Our algorithm takes this into account by trying several different subsets of the flow requirements. It selects that subset which is routed along the 'shortest paths' when the real cost and estimated cost are very nearly equal.

After each iteration the lower bounds y_m must be updated. The subset of flow requirements routed at any given iteration may equal the null set which means that the estimated costs expressed by the y_m's are too optimistic in that the real costs are significantly greater for each link. If this happens then the lower bounds or weights y_m are marginally increased in order that they become more realistic.

This algorithm has been implemented in a computer program which can generate several solutions to the problem of Equations (1) through (6) by varying some built in parameters. These parameters determine how well the estimated costs (y_m) approximate the real costs. They also control the magnitude of the increase of the y_m's when no flow was routed during the previous iteration.

An additional feature that has been built into the program is the capability of improving a solution obtained by a human designer or from the above algorithm. The program compares all the links included in the solution and determines the link through which traffic has been routed at the maximum cost per unit of flow. This maximum cost may of course apply to only a portion of the flow through the link because of the nature of the link costs. In any case, the program places a restriction on the flow through this link and the above algorithm is then used to find a solution.

No provisions for constraining the number of earth stations were included in the above discussion. Suitable methods are best devised by considering the nature of each specific network problem. For instance, in the Canadian situation, most traffic centres lie close to the southern border in a pattern which approximates a straight line. Two promising options in such a case are as follows:

1. Carry out an initial unconstrained optimization first and then progressively eliminate, possibly by human intervention, those earth stations carrying the least traffic.

2. For the n-earth station case, simply constrain the solution to the first n stations introduced by the algorithm.

For more general cases techniques such as those described by Mirowsky[6] would be useful.

3.3 EXPERIENCE

Experience to date in using the algorithm indicate that it is capable of generating solutions whose costs are within 0 - 10 per cent of the optimum solution as generated by linear programming methods at about one per cent of the cost of such methods. The algorithm has been applied to networks with up to 60 nodes, 175 links and 100 commodities.

4. SUMMARY

The cost-demand strategy described herein makes it possible to divide the overall task of designing a satellite/terrestrial trunk-communications system into two distinct tasks of similar difficulty. We have shown that the two tasks can be carried out separately until close to the end of the analysis. However, the final result is the determination of the most cost-effective satellite/terrestrial trunk network. This network is specified by the intersection point on the cost-demand surface. We expect that close interaction between those doing the two tasks could reduce the 'area' on Figure (5) to be 'searched' for an optimum system.

We have described a heuristic algorithm for finding network solutions necessary for determining demand curves. Results so far indicate that it is capable of efficiently generating good sub-optimal solutions.

5. REFERENCES

1. Formby, J.A. *Computer procedure for bounding the chromatic number of a graph. Combinatorial Mathematics And Its Applications.* (Edited by D.J.S. Welsh), London, Academic Press Inc., 1969, pp. 111-114.

2. Frank, H. and W. Chou. *Topological optimization of computer networks.* Proceedings of the IEEE, November 1972, pp. 1385-1397.

3. Fratta, L., Gerla, M. and L. Kleinrock. *The flow deviation method: an approach to store-and-forward communication network design.* Networks, Vol. 3, No. 2, 1973, pp. 97-133.

4. Jewell, W.S. *Multi-commodity network solutions.* ORC 66-23, Operations Research Centre, University of California, Berkeley, 1966.

5. Kleitman, D.J. and A. Claus. *A large scale multicommodity flow problem: Telpak.* Proc. Symposium Computer-Communications Networks and Teletraffic, New York, Polytechnic Press, 1972, pp. 335-338.

6. Mirowsky, B.J. *Designing a system for maximum return.* Presented at the IEEE Electronics and Aerospace Systems Conference, September 1973.

7. Peck, J.E.L. and M.R. Williams. Algorithm 286, Exam Scheduling, Comm. ACM. Vol. 9, No. 6, 1966, pp. 433-434.

8. Rothfarb, B. and M. Goldstein. *The one-terminal telpak problem.* Operations Research, Vol. 19, No. 1, 1971, pp. 156-169.

9. Welsh, D.J.A. and M.B. Powell. *An upper bound for the chromatic number of a graph and its application to timetabling problems.* Computer Journal, Vol. 10, 1967-68, pp. 85-86.

10. White, W.W. *Mathematical programming, multicommodity flows, and communications nets.* Proc. Symposium Computer-Communications Networks and Teletraffic, New York, Polytechnic Press, 1972, pp. 335-338.

11. Williams, M.R. *The coloring of very large graphs.* Proc. of the Calgary International Conference on Combinatorial Structures and Their Application, Gordon and Breach, Science Publishers, 1969, pp. 477-478,

12. Zangwill, W.I. *Nonlinear programming.* Englewood Cliffs, N.J., Prentice-Hall, 1969.

CRC DOCUMENT CONTROL DATA

1. ORIGINATOR: Communications Research Centre,
Department of Communications,
Ottawa, Canada.

2. DOCUMENT NO: CRC Technical Note No. 663

3. DOCUMENT DATE: January 1974

4. DOCUMENT TITLE: The Planning of Domestic Communication Networks.

5. AUTHOR(s): G.A. Neufeld, R.R. Bowen and A.R. Kaye

6. KEYWORDS: (1) Planning
(2) Communications
(3) Networks

7. SUBJECT CATEGORY (FIELD & GROUP: COSATI)

Mathematical Sciences

Operations research

12-02

8. ABSTRACT:

Techniques for the synthesis of domestic long-haul communication systems to meet forecast traffic requirements must provide for cost-effective integration of satellites with existing and new terrestrial systems. A methodology that may be used to carry out this synthesis is described. One of its important features is the separate investigation of the characteristics of satellite and terrestrial systems with merging of the results towards the end of the study. This allows activities to be carried out concurrently rather than consecutively so that both sets of results may be more up-to-date and so more useful when they are merged. Optimization algorithms are described which are of general interest in that they can be used to optimize any network having step-like link costs.

9. CITATION: ______________________________

Communications
Canada