A Simulation Study of Routing and Control in Communications Networks*

By J. H. WEBER

(Manuscript received March 23, 1964)

A set of studies has been undertaken to develop guidelines for the design and operation of communications networks with automatic alternate routing. Comparisons are made of engineered costs and overload capability of networks using several alternate routing configurations, and employing a number of different operating and control procedures. The traffic model selected consists of a 34-node network abstracted from the U.S. telephone toll network, with basic load levels obtained from field data. The overload evaluations were made using a simulation program prepared for the IBM 7090 computer.

I. INTRODUCTION

In a recent paper[1] the results of some preliminary comparisons of two alternate routing configurations for communications networks were reported. Those results indicated that for small networks (six or fewer nodes) with low traffic densities a symmetrical or unrestricted routing pattern is superior to a hierarchy similar to that in use in the U.S. toll network, while for higher traffic densities there appeared to be little difference in the network behavior in terms of economy and reaction to overloads.

Subsequently, a new simulation program has been constructed[2] and substantially larger networks have been examined to provide a more meaningful guide to network design under various circumstances of geography and load level. An additional configuration, called the "gateway," as well as several operating and control variations, has been examined. The latter include stage-by-stage operation with and without crankback (return of routing control to a previous node for rerouting when blocking is encountered at an intermediate switching point), limitation of number of links per call in symmetrical networks, and trunk reservation for first-routed traffic only.

* Presented at the Fourth International Teletraffic Conference, London, July, 1964.

The results show that:

(1) There is little difference in network cost or overload capability between hierarchical and symmetrical networks at the load densities considered.

(2) The gateway network (a two-level hierarchy with no interregional high-usage groups) requires substantially more trunking and switching than either the hierarchical or symmetrical networks and shows no significant difference in overload performance.

(3) Restriction of alternate routing in symmetrical networks improves performance at all levels of load.

(4) The use of crankback is a disadvantage for symmetrical networks with a high traffic density, at all levels of overload. It offers a slight advantage for symmetrical networks with lower traffic intensities and does not appear to have any significant effect on the performance of hierarchical networks.

(5) Trunk reservation for first-routed traffic on a dynamic basis improves the performance of almost all networks examined, for all load conditions, and displays no detrimental effects.

II. THE SIMULATION

The simulation program used in these studies is described in Ref. 2. It has many of the capabilities of the program described in Ref. 1, but has been reprogrammed to accept networks with heavier loads and to operate more efficiently. A number of additional features have also been provided.

The program is basically capable of simulating networks of up to 63 nodes, with arbitrary alternate routing patterns and stage-by-stage call forwarding. There is no congestion or delay allowed at switching points, all congestion being assumed due to trunk shortages. Calls which fail to complete initially may be abandoned with a fixed probability or retried after a constant or exponentially distributed interval. Any prespecified number of trunks can be reserved for first-routed traffic only, and calls may "crank back" or return to a prior node if blocked at some point in the network. The maximum-size network which can be accommodated is largely determined by the number of simultaneous calls in progress, which may have a maximum of about 6000. Traffic loads are specified on a point-to-point basis, with arbitrary proportions in each direction, and may be changed linearly at any time during the run. That is, mean arrival rates can change linearly in time during the run at any rate and between any bounds. (Another modification of the program allows the use of two priorities of traffic and mixtures of direct and store-and-forward traffic, with trunk reservation by traffic type and priority. This version was not used for the studies described herein, however.)

In order to accommodate larger networks more efficiently, the program was written in several sections. The first of these accepts the basic load inputs (mean point-to-point loads and holding times) and generates call arrival times and holding times, which are then stored on magnetic tape. This tape is then used as input to the simulation program, which processes the calls through the simulated system and prints out raw data on trunk utilizations and call histories on two magnetic tapes. These tapes are presented to the output processor programs, which provide the appropriate reduced outputs.

For convenience in preparing the input data, the main section of the program has been arranged to determine its own routing for symmetrical and hierarchical networks, given the numbers of trunks and the distances for symmetrical networks, or the homing arrangements for hierarchical networks.

The output statistics are reported at prespecified time intervals, and these subinterval results may then be used as samples for a final output containing both means and standard deviations of all relevant quantities. The quantities which are printed out are as follows:

(1) point-to-point traffic loads at the end of the run (input data).

(2) routing tables for all point-to-point traffic items.

(3) means and standard deviations of the following measured quantities for each point-to-point traffic item:

(a) blocking probability

(b) average delay and distribution of delays for retried calls

(c) average number and distribution of number of links per call.

(4) means and standard deviations of the following measured quantities for each trunk group (obtained by switch count measurements):

(a) number of trunks present in each group (input data)

(b) number of trunks reserved for first-routed traffic in each group (input data)

(c) total carried load in erlangs on each group (and per cent occupancy)

(d) first-routed carried load on each group (and per cent of total)

(e) alternate routed carried load on each group (and per cent of total).

(5) means and standard deviations of measured over-all network quantities as follows:

(a) over-all average blocking probability, $\bar{B}$, given by

$$\bar{B} = \frac{\sum_{ij} a_{ij} B_{ij}}{\sum_{ij} a_{ij}}$$

where

a_{ij} = offered load between nodes i and j, and
B_{ij} = blocking probability of calls offered between nodes i and j

(b) average number and distribution of number of links per call
(c) weighted average delay and delay distribution for retried calls
(d) total number of trunks in the network (input data)
(e) total trunks reserved for first-routed traffic (input data)
(f) total carried load and over-all occupancy
(g) total carried first-routed load
(h) total carried alternate routed load.

(6) the "space dispersion," D, of the blocking probability, delay distribution and links per call distribution, given by

$$D_B = [(\sum_{ij} a_{ij} B_{ij}^2 / \sum_{ij} a_{ij}) - \bar{B}^2]^{\frac{1}{2}}$$

for the blocking probability, and similar expressions for the other quantities. This serves as a measure of the variation in grade of service provided to various traffic items in the network depending upon their origin and destination.

(7) the following input parameters:

(a) maximum number of links allowed per call

(b) number of stages of crankback allowed (That is, the number of steps a call is allowed to back up before progressing forward after having reached a point of congestion.)

(c) percentage of calls to be retried
(d) retrial time distribution and mean value
(e) holding time distribution and mean value
(f) number of nodes
(g) number of reporting intervals and their lengths
(h) number of reporting intervals to be collected for final processing
(i) routing pattern (hierarchical or symmetrical)
(j) interval between switch counts for determination of trunk information.

The simulation runs quite rapidly, processing about 500,000 calls per hour using the IBM 7090 computer for a 34-node network with a total load of about 5000 erlangs, excluding the traffic generation. If traffic generation time is included, the processing rate drops to about 375,000 calls per hour. If several networks are evaluated using the same traffic input, as was done in these studies, however, the traffic need be generated only once, and the same tape can be used as input to the simulation any number of times. Several dozen simulation experiments were made for the studies described below, but only eight traffic tapes were generated.

III. GENERAL PROCEDURE

The evaluation procedure encompasses the following steps:

(1) select a geographical area, including switching center locations;

(2) select a basic traffic model on which to base network engineering;

(3) engineer networks to a given grade of service using each of the routing procedures to be considered;

(4) determine the costs of each of the networks so engineered;

(5) change the loads to correspond to reasonable patterns of overload or shifting load;

(6) using simulation, measure the performance of each of the networks under the load changes used in step (5);

(7) repeat steps (5) and (6) for each of the control and operating variants considered.

These steps will be described in detail in the following sections.

3.1 *The Geographical Region*

Although it is not possible to select a geographical region (or regions) which will be typical of all situations, it is desirable to find an area which at least has the capability of accommodating a sufficient number of nodes to adequately exercise the various routing patterns to be examined and of reacting to realistic load fluctuations. The region should also contain both densely and sparsely populated sectors, which to some extent must exist in all real networks. (A uniform or arbitrarily variable traffic distribution would probably not be a valid test, since actual telephone traffic varies with the population density, higher-density areas having large amounts of traffic within and among them, and sparsely populated areas being lower in traffic to all destinations.) Since the geographical region will utimately require a traffic model to be superimposed upon it, an area for which actual traffic data is obtainable is again more likely to represent reality than one for which traffic quantities need to be invented.

A single region which appears to meet most of the criteria specified above exists on the West Coast of the U.S.A. The states of California, Washington, Oregon and Nevada are almost entirely administered by two local telephone operating companies for toll purposes and represent a region which ranges from sparsely populated areas such as Nevada and eastern Oregon to sections such as Southern California, which contains the Los Angeles and San Diego metropolitan areas. Fig. 1 is a map of this region, showing the 34 switching centers used. Although there are many more than 34 toll switching offices in this region, only the control switching points (CSP's), which make up the offices of the top three levels of the U.S. toll network hierarchy,[3] are included. (Las Vegas, though actually a toll center, is assumed to be a primary center.) All traffic which both originates and terminates in the region, however, is included, as will be discussed below.

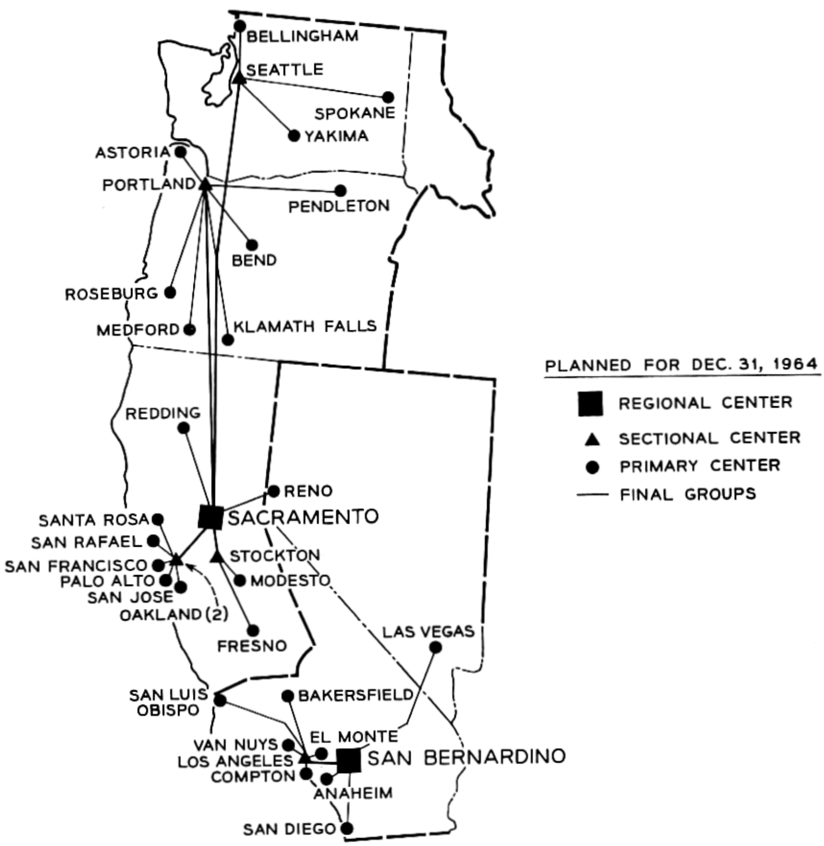

Fig. 1 — U.S. West Coast traffic region.

3.2 *The Basic Traffic Model*

The basic traffic model used for engineering the various networks was developed from actual message records of the Pacific Telephone and Telegraph Company and the Pacific Northwest Bell Telephone Company. These records include the total messages for a period of ten consecutive business days during June 1962. They provide total messages and message minutes from every toll switching center in the area to every other. Traffic which originates or terminates in other than the four-state area is not considered, nor is traffic which is not carried on the toll (or long distance) network. Traffic originating and/or terminating at offices of connecting companies, but which is carried on the toll network, is included.

In order to obtain a busy hour traffic base from the ten-day records, it was assumed that ten per cent of the total traffic was offered during each day and ten per cent of the day's traffic was presented during the busy hour. Therefore the busy hour traffic load was assumed to be one per cent of the total ten-day message load.

Traffic between toll centers of the fourth rank, which are not explicitly included in the 34-node model, is handled in two ways, giving rise to two networks with different traffic densities. In the first of these, called the "full-traffic" or "full-load" network, all traffic between toll centers is added to that between the centers on which they home. Traffic between toll centers homing on the same control switching point is eliminated. For example, referring to Fig. 2, traffic between toll centers A and B is added to the traffic between control switching points D and E, as is the traffic between A and E, and between B and D. Traffic between toll centers B and C and between points A and D, and B and E is eliminated.

In the second network, called the "reduced-traffic" or "reduced-load"

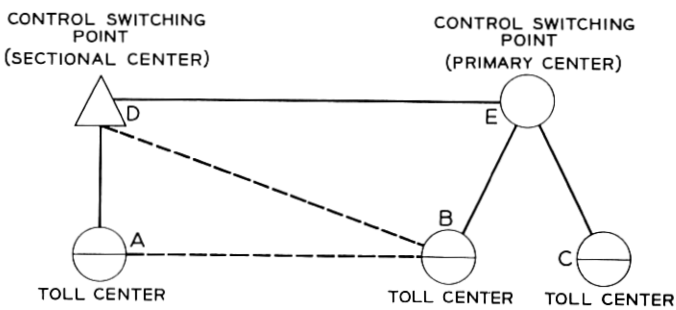

Fig. 2 — Disposition of toll center traffic.

network, it is assumed that some of the traffic between toll centers is carried on high-usage groups, and only the overflow is carried on the groups between the CSP's. Therefore traffic between toll centers such as A and B, in which there may be two routes before the CSP-CSP trunk group is reached, is assumed to overflow only 10 per cent to the D-E group. Traffic between a toll center and a CSP, such as that between A and E, may have only one route before the CSP-CSP route is reached, and 20 per cent of this traffic is then assumed to be offered to route D-E. Traffic between toll centers such as B-C, or between toll centers and the CSP on which they home, such as A-D, is again eliminated.

The net effect of these assumptions is to develop a total network load of 4764 erlangs for the full-load case, and 2031 erlangs for the reduced-load network. The maximum point-to-point load for the full traffic network is 158 erlangs, and the maximum load originating and terminating at any node is 848 erlangs. The minimum point-to-point load is 0.01 erlang, and the smallest node has 26 erlangs originating and terminating at it. For the reduced network the maximum point-to-point load is 84 erlangs and the minimum is zero. The total traffic originating and terminating at the largest node is 288 erlangs, and at the smallest is 15 erlangs. A tabulation of the total loads originating and terminating at each point in both networks is given in Table I.

3.3 *The Network Configurations*

Five specific networks of three configuration classes for the full-traffic model and two networks in two classes for the reduced-traffic model were examined. The first class of networks is hierarchical in structure, similar to that in use in the Bell System toll network. In these networks, trunk groups are defined as high-usage, which may overflow traffic to alternate routes, or final, which may not. The apportionment of trunks among high-usage and final routes is decided on an economic basis.[4] Both two- and three-level hierarchies were examined in the full traffic model, while only two levels were used for the reduced traffic case. The routing for these networks is shown in Fig. 3.

In Fig. 3(a) the basic elements of a two-level hierarchy are shown. Calls from node 1 to node 2 will, if unable to use the direct route, attempt to reach node 4, from which the only allowable choice is the final route 4-2. If unable to reach node 4, calls will then attempt to reach node 3, from which they will attempt the direct route 3-2, finally overflowing to the final route 3-4. Calls from 2 to 1 will reverse the procedure, attempting to reach node 3 and overflowing to node 4. Calls initially routed

Table I — Switching Center Loads

Switching Center	Total Originating and Terminating Traffic In Erlangs	
	Full Load	Reduced Load
Bellingham, Wash.	94.17	34.81
Seattle, Wash.	533.25	272.59
Spokane, Wash.	141.89	71.96
Yakima, Wash.	175.17	59.73
Astoria, Oregon	26.92	16.87
Bend, Oregon	26.65	15.19
Klamath Falls, Oregon	33.12	19.60
Medford, Oregon	62.42	37.08
Pendleton, Oregon	51.35	21.11
Portland, Oregon	468.05	233.56
Roseburg, Oregon	37.37	17.85
Las Vegas, Nevada	116.74	82.04
Reno, Nevada	138.21	69.44
Fresno, Calif.	306.11	140.81
Modesto, Calif.	132.55	81.86
Stockton, Calif.	206.21	105.38
Redding, Calif.	87.94	30.88
Sacramento, Calif.	539.32	173.63
San Jose, Calif.	459.24	188.90
Oakland 4M, Calif.	848.22	288.14
Oakland Fr., Calif.	524.25	194.06
Palo Alto, Calif.	251.14	108.99
San Francisco, Calif.	375.40	198.45
San Rafael, Calif.	134.12	40.26
Santa Rosa, Calif.	386.08	107.64
Bakersfield, Calif.	167.84	93.56
San Luis Obispo, Calif.	109.15	45.59
Compton, Calif.	447.04	190.15
Los Angeles, Calif.	699.27	242.63
El Monte, Calif.	361.88	179.54
Van Nuys, Calif.	426.75	174.32
Anaheim, Calif.	158.45	73.78
San Bernardino, Calif.	576.81	202.72
San Diego, Calif.	424.33	248.71
Total (orig. plus term.)	9527.41	4061.54

along the final route chains, such as those from 1 to 3, have only a single choice of route.

The three-level network, shown in Fig. 3(b), allows a somewhat more complicated routing pattern. Calls from 1 to 2 in this network will attempt to reach nodes 2, 4, 6, 5 and 3 in that order, and all other routes will follow a similar pattern of hunting from low to high level in the distant region, and from high to low level in the home region. In no event can a call use more than one interregional trunk, and calls always travel up the hierarchy in the home region and down in the distant region. A fuller description of the process is given in Ref. 3. This restrictive routing

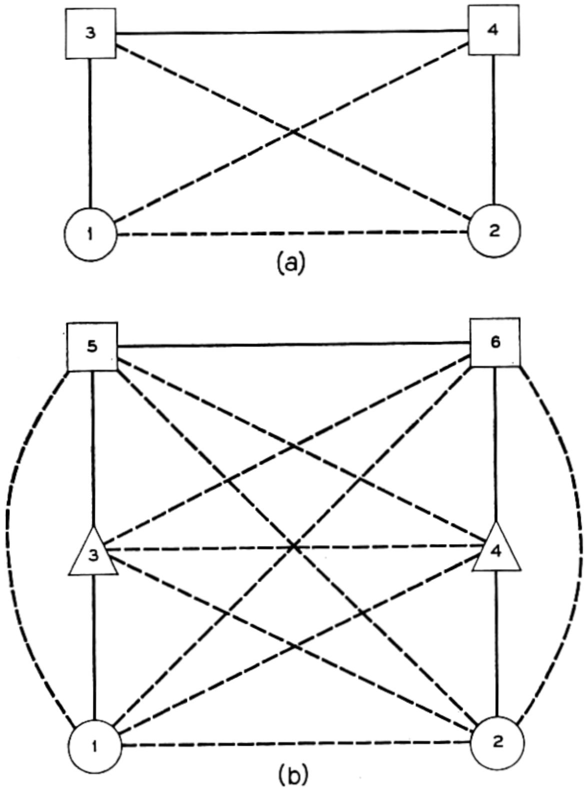

Fig. 3 — Organization of hierarchical networks.

pattern allows alternate routing to proceed without fear of "ring-around-the-rosie" or "shuttling" (which are types of looping routes), even though no information is carried with the call other than its destination code. The two-level hierarchical networks actually used contained six higher-level offices, located at Seattle, Portland, Sacramento, Oakland, San Bernardino, and Los Angeles. The three-level network took Portland, Sacramento, and San Bernardino as highest-level, or regional, centers; leaving Seattle, Oakland, and Los Angeles as middle-level, or sectional, centers. A sketch of the Washington-Oregon section of the full-load, two-level hierarchy is shown in Fig. 4.

The other network configuration examined for both load levels is the symmetrical network, in which alternate routes are selected approximately according to their total length. In all such networks studied, trunks are arbitrarily eliminated on links with less than 2 erlangs of directly offered traffic, and routing is then established using equipped links. Fig. 5 shows the trunk group layout for the Washington-Oregon section of the full-load symmetrical network. A basic restriction in all networks is that at most five outgoing choices are allowed from any node to any other, it being considered that further choices would lead to excessively circuitous routes. In addition, no route is allowed which is more than 1.5 times as long as the shortest nondirect route, or exceeds the shortest nondirect route by more than 2 links. These numbers were arrived at by trial and error and produced the most economical network for the full-load case, although they were not very critical in the determination of network cost or capability. Two symmetrical networks are studied in the full-load case, one which matches the blocking performance of the other networks at engineered loads, and one which has a higher blocking, as described below. Only one symmetrical network is used for the reduced-load model.

The method by which routes are selected is as follows. Initially, the shortest route between each two points is found. The route to the nearest neighbor node on this route is then listed as the first-choice route. The link from the originating node to the nearest neighbor node along the first-choice route is then made ineligible, and the shortest route again found. The link to the nearest neighbor node along this route is then

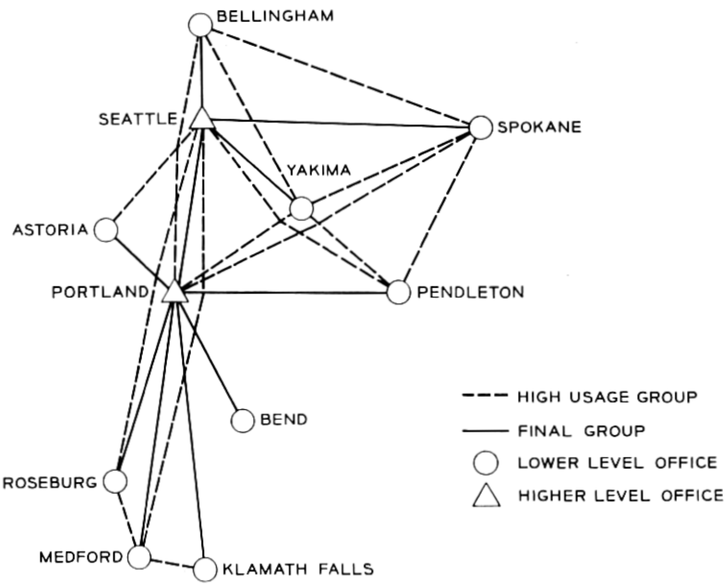

Fig. 4 — Full-load, two-level hierarchy — Washington-Oregon.

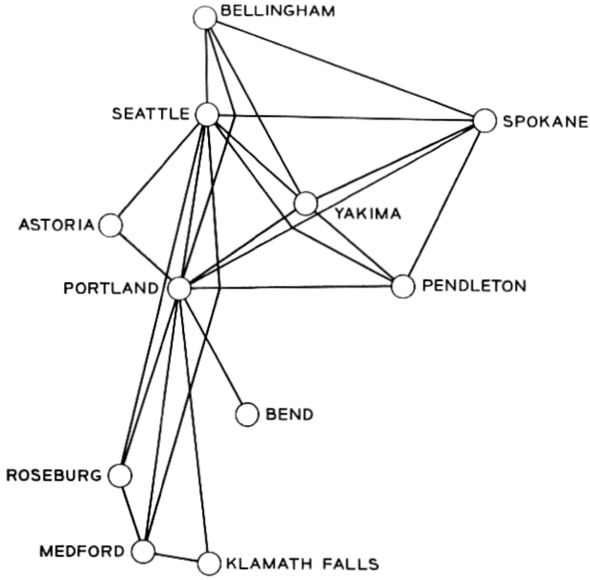

Fig. 5 — Full-load symmetrical network — Washington-Oregon.

denoted the second-choice route, and the distance and number of links calculated and compared with the first nondirect route. The entire procedure is repeated until no route falls within the distance ratio and link difference criteria, or five routes are selected, whichever occurs first. At this point the process is terminated and the routing table established. For example, in Fig. 5, to go from Yakima to Medford, the first-choice route is via Portland, the second is via Seattle, and the third is via Pendleton.

The third network configuration, considered in the full traffic case only, is the gateway network. This is essentially a two-level hierarchy with the interregional high-usage groups removed, as shown in Fig. 6 for Washington and Oregon. Traffic and trunks are therefore concentrated along the access routes to the gateway switching center and on the interregional finals. Although this kind of system clearly requires more trunks and trunk miles than a hierarchy to carry the same loads, it has been conjectured that savings in line and terminal equipment could be effected because of the large trunk cross sections involved. It also has been thought that this scheme might provide improved performance under shifting loads, which hypothesis is examined in this study. The gateway network studied assumed the gateway switches to be located at the same points as the higher-level offices in the two-level hierarchical networks.

In all of the above networks, stage-by-stage routing similar to that in the U.S. toll network is used. That is, once a call has reached a certain point in its path, its route selection is independent of its past history, and it is unable to back up and find another route out of a prior node. (This is not true if crankback is allowed, as will be discussed later.) In the symmetrical networks, the previous route is considered to the extent of preventing a call from returning to a node through which it has already been switched. In the hierarchy and gateway, this restriction is implicitly provided by the logic of the routing structure.

In sum, the networks examined are as follows:

(1) Full-load model

- (a) two-level hierarchy
- (b) three-level hierarchy
- (c) symmetrical
- (d) symmetrical with high blocking
- (e) gateway.

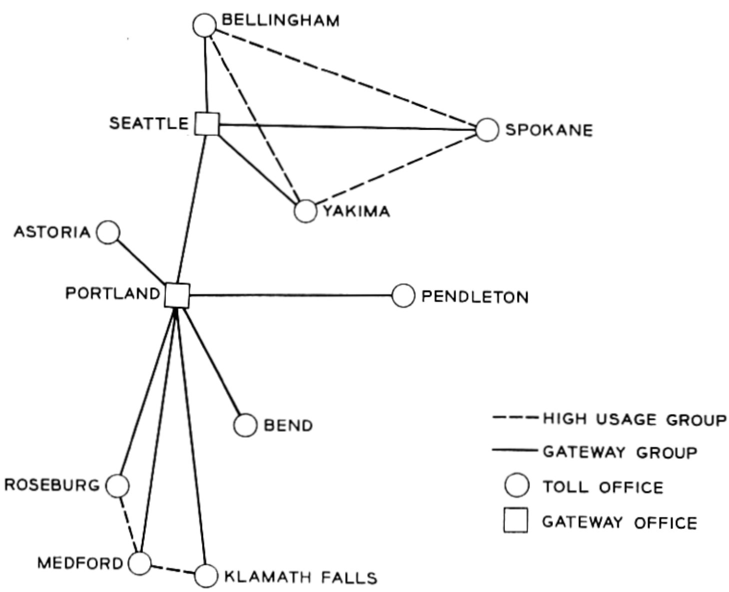

Fig. 6 — Gateway network — Washington-Oregon.

(2) Reduced-load model

 (a) two-level hierarchy
 (b) symmetrical.

3.4 *Engineering Procedures*

The size and complexity of the networks considered are such that manual engineering procedures or trial and error methods are not feasible. Accordingly, computer programs were prepared which established at least an initial network, which could then be adjusted if required by trial and error using the simulator. The objective for all networks was to attain an over-all average blocking probability of 0.01, with as small a dispersion of individual point-to-point probabilities as possible. This is a somewhat different criterion than the one normally used in existing hierarchical alternate routing systems, which specify the blocking probability observed on the final route, but it is closer in philosophy to local systems and others in which blocking probabilities produced by the system are the same to all customers.

The hierarchical networks were engineered with the aid of a computer program which essentially follows the procedure outlined in Ref. 4. Using this method, traffic is transferred from the direct to the alternate route when the direct route becomes so inefficient that the cost of adding a trunk to it is more than the cost of carrying the traffic on the alternate route. No account was taken of the nonrandomness of overflow traffic[5] or of the nonindependence of different links in the network. The errors resulting from these assumptions were not large and were corrected where required by trial and error using the simulator.

The process for engineering a symmetrical network is less well developed, and no method for designing an optimal, or even necessarily a very good network, exists. However, a program is in existence[6] which is capable of designing networks which will closely meet a desired blocking probability, using prespecified routes which are fully determined from origin to destination. It was necessary, in order to use this program, to convert the shortest route procedure described in Section 3.3 above to one which provides the full route rather than simply the order of hunt over the adjacent nodes. This resulted in networks which were engineered using a slightly different routing arrangement than the simulator actually used, and this, in conjunction with the basic assumptions implicit in the engineering program of random overflow traffic and independent links, led to blocking probabilities in the final network which were somewhat higher than desired. These were corrected for the purposes of comparing network configurations, but for certain studies of various methods of operating a symmetrical network, the networks with high blocking were retained.

3.5 *Load Changes*

Three patterns of load changes were used to measure the performance of the various networks under shifting load conditions. The two load changes examined in Ref. 1, uniform overload and overload of all traffic to and from a particular node, were avoided because of their limitations. The first case, uniform overload, represents a situation which is thought not to ordinarily occur in large real systems, and both load models are likely to obscure differences in behavior of competing networks, since the networks tend to be completely saturated or are limited by the specific overloaded nodes. Instead, three patterns of shifting loads, in which the total offered network load remained approximately unchanged, were used.

The first of these, called the "Christmas load," represents a type of shifting load normally seen in the U.S. on Christmas Day and on a few other special occasions. On these days, the normal long distance business traffic disappears and is replaced by a large volume of residential traffic. Typically, the increased traffic is of substantially longer haul than is the normal day traffic, so the phenomenon observed is that short-haul traffic decreases, but long-haul traffic increases. In order to represent this in the sample Pacific network, the network was broken down into four areas, consisting of Washington, Oregon, Northern California and Southern California. (Northern Nevada was included with Northern California and Southern Nevada with Southern California.) All intra-area traffic was reduced to 60 per cent of its normal value, and inter-area traffic was increased to from 150 to 275 per cent of its normal value, depending upon the distance. The total network load was 94 per cent of its normal value, as shown in Table II. Although these changes may appear extreme, they are not thought to be out of line with what actually occurs in the U.S. on Christmas and were applied to both full and reduced traffic models.

The second load change examined is not typical of any actual situation, but was designed to evaluate the effectiveness of the various networks in shifting load from an overloaded trunk group to a simultaneously underloaded one. In order to do this for the full-load network, all traffic items originating or terminating at the Oakland 4M machine, the largest office in the network (33 traffic items, total load 848 erlangs

TABLE II — CHRISTMAS LOAD CHANGES

Traffic	% of Normal Day Busy Hour Load
Intraregional	60
Washington-Oregon	150
Washington-Northern Calif.	210
Washington-Southern Calif.	275
Oregon-Northern Calif.	175
Oregon-Southern Calif.	230
Northern Calif.-Southern Calif.	150
Network	94

representing about 20 per cent of the network load), were either halved or doubled at random, although this was slightly modified so that the total network load remained at 99.5 per cent of its normal value. The normal and modified values of the Oakland loads are shown in Table III. In the reduced traffic model no single office had enough traffic to cause substantial changes in total network performance, so the halving and doubling were done at Seattle, Oakland and Los Angeles, which have total loads of 767 erlangs, representing about 40 per cent of the total network load. In this case the total network load increased by about 8 per cent. It is to be emphasized that this set of loads does not represent any expected realistic situation, but is a completely artificial test of the effectiveness of automatic rerouting under most favorable conditions.

The third load change examined was actually a series of load changes based on an assumed movement of the busy hour from north to south during a four-hour period. It was further assumed that, relative to the busy hour, an area's load was reduced 5 per cent in the adjacent hour, 10 per cent in the second hour, and 20 per cent in the third. Traffic between two areas (defined in the same way as for the Christmas loads) was taken to be the arithmetic mean of the levels of the terminal offices. That is, a traffic item between an area which is in its busy hour and one two hours distant is assumed to be reduced 5 per cent from its busy hour value.

Since the networks were engineered based on a single over-all load value, some normalization was done so that the over-all network load remained approximately constant. There were also some limitations in the simulation which prevented the desired changes from being reached exactly, but the final loads were quite close to the desired value. The sequence of changes is shown in Fig. 7. The first part of the line represents the basic engineered load, followed by the changes in each area's traffic as shown. The ordinate is a relative scale, so all loads are given as multiples of the basic value. The ramps between the hours were actually simulated as shown, but no measurements were taken during these periods. The inter-area traffic levels are not shown, but are arithmetic means of the levels of the terminal nodes, as described above.

This load change, which was applied to both full- and reduced-load networks, is designed to analyze a situation similar to that in the entire U.S.A., which has several time zones with a different busy hour in each. Although such differences are, of course, not actually observable in the network selected, which runs essentially north and south, we can, for the purposes of modeling, assume it runs from east to west, and is expanded in its dimensions. In this case, time zone changes like those postulated for this "busy hour load" would in fact be observed.

TABLE III — OAKLAND VARIATION LOADS (FULL-LOAD NETWORKS)

Traffic between Oakland 4M and	Traffic Loads in Erlangs	
	Normal Loads	Changed Loads
Bellingham, Wash.	3.58	7.16
Seattle, Wash.	26.64	13.32
Spokane, Wash.	3.31	6.62
Yakima, Wash.	2.39	1.20
Astoria, Oregon	0.23	0.46
Bend, Oregon	0.29	0.15
Klamath Falls, Oregon	1.08	2.16
Medford, Oregon	1.77	0.88
Pendleton, Oregon	0.50	1.00
Portland, Oregon	23.77	11.88
Roseburg, Oregon	0.50	1.00
Las Vegas, Nevada	3.78	1.89
Reno, Nevada	15.98	31.96
Fresno, Calif.	30.49	15.24
Modesto, Calif.	12.59	25.18
Stockton, Calif.	25.86	12.93
Redding, Calif.	8.68	17.36
Sacramento, Calif.	94.23	47.12
San Jose, Calif.	76.79	153.58
Oakland Fr., Calif.	122.95	61.47
Palo Alto, Calif.	45.47	90.94
San Francisco, Calif.	45.12	22.56
San Rafael, Calif.	26.18	52.36
Santa Rosa, Calif.	86.80	43.40
Bakersfield, Calif.	7.32	14.64
San Luis Obispo, Calif.	3.48	1.74
Compton, Calif.	33.57	67.14
Los Angeles, Calif.	60.73	30.36
El Monte, Calif.	26.74	53.48
Van Nuys, Calif.	28.77	14.38
Anaheim, Calif.	8.16	16.32
San Bernardino, Calif.	8.65	4.32
San Diego, Calif.	11.81	23.62
Total	848.21	847.82

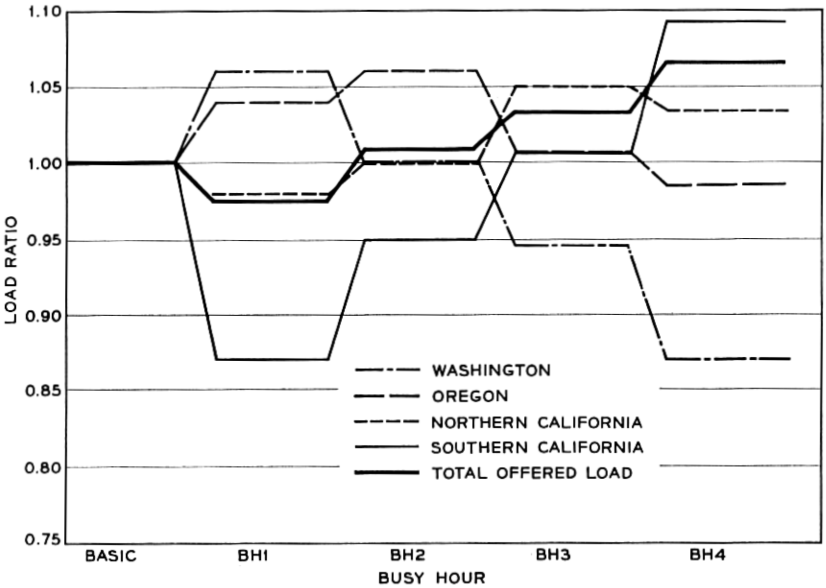

Fig. 7 — Busy hour load changes.

3.6 *Evaluation Criteria*

A communications network which must be engineered to meet a specific set of demands for service without being excessively costly, and which will then be subjected to demands for which it was never designed, is not easily evaluated by a single figure of merit, or even by a small number of parameters. The weight of overload performance versus engineered economy, performance under overload A as opposed to that under overload B, and service to traffic between points i and j as opposed to that provided between points k and m, provide ample opportunity for conflicting requirements. This is, of course, in addition to nontraffic considerations such as survivability, ease of engineering, administration and control, or ability to provide other services such as data and private line.

Nevertheless, in order to make a comparative evaluation of various network configurations, a set of criteria must be adopted which can be evaluated for each network under study and which will reflect the basic considerations of cost and service quality under all conditions.

The criteria which have been selected are four in number, two relating to cost and two relating to grade of service. They are:

(1) number of trunks required to provide the desired grade of service [of $P(0.01)$] at engineered load

(2) number of trunk miles required to provide the desired grade of service at engineered load

(3) the over-all weighted average blocking, $\bar{B}$, as defined in Section II above

(4) the dispersion of blocking D_B (subsequently denoted simply D), as defined in Section II above.

Items (1) and (2) above can be provided with costs to derive approximate network costs, which will vary depending upon the cost of switching, terminal equipment and line facilities. This has been done for a few typical costs. The costs so derived are approximate because the trunk miles are defined as point-to-point airline miles, which is not the way actual facilities would normally be routed.

Items (3) and (4) are measures of service quality. $\bar{B}$ is by itself a measure of over-all network performance, and it is directly related to carried load. However, there may be severe distortions in the point-to-point blockings which would yield a low $\bar{B}$ but might still leave certain customers with extremely poor service. The inclusion of D as a criterion will help to identify such a situation and ensure that network service is evaluated on a basis of balance as well as blocking level.

3.7 *Operating and Control Procedures*

The variations in operating procedures and the control methods employed all have the effect of changing the amount of alternate routing, normally also making different numbers of routes available to various point-to-point traffic items. Two control procedures were investigated for both the hierarchical and symmetrical networks. These were one-stage crankback and trunk reservation for first-routed traffic only (subsequently referred to simply as trunk reservation). One-stage crankback allows a call which has reached a point from which it is unable to proceed to back up one link along its previous route and attempt to complete via another route. This has been proposed as both a traffic improvement measure and as a means for allowing machine troubles to be circumvented without customer retrials. The investigation here relates, of course, only to its effect on the traffic capacity of the network. Trunk reservation allows only first-routed traffic to seize the last idle trunk in a group. Alternate routed calls can be served only if at least $m + 1$ trunks are idle, where m is the number of trunks reserved. This procedure tends to maximize the number of calls which are carried on direct links at the expense of those carried on alternate routes. It also reduces group efficiencies somewhat, and the question is whether the reduction in circuitous routing is enough to compensate for this.

Finally, for symmetrical networks only, the maximum number of links per call was varied. In the case called "full routing," a maximum of five links was allowed for any call in the network. In the case called "limited routing," a maximum of only three or four links per call was allowed, depending upon the connectivity between the originating and terminating points of the call. This restriction, of course, reduces the average number of links per call, at the same time reducing the number of routes possible between any two points.

IV. ANALYSIS OF NETWORK CONFIGURATIONS

4.1 *Facility Requirements*

It is difficult to arrive at an accurate measure of the cost differential between the various network configurations, since costs of trunk terminations, switching, and trunk lines vary from place to place and from network to network. However, it is expected that the relative costs of the various network configurations can be deduced from the number of trunks and the number of trunk miles by applying appropriate factors related to the distribution of trunk lengths and the types of switching and transmission equipment in general use in any given situation. If the unit costs of switching equipment, or of control features inherent in the routing plan, are significantly different for different networks, the magnitudes of these differences can be balanced against the differences in trunks and trunk miles to again deduce the total network relative costs. It should also be noted that the distances used for the trunk length calculations are based on airline mileage between originating and terminating points, which is ordinarily somewhat shorter than actual facility route mileage. This discrepancy can be corrected by introducing multiplying factors when determining network costs for any actual case.

Table IV shows the number of trunks and trunk miles required to provide the noted grade of service for each of the networks under consideration, both in absolute value and as per cent difference from the two-level hierarchy, which was arbitrarily selected as the standard. Although the blocking probabilities are not exactly the same for all networks due to inaccuracies in the engineering procedures and statistical fluctuations in the simulations, they are quite close.

The differences in facilities required for the various networks, with the exception of the gateway, are quite small, amounting to at most 4.1

TABLE IV — COMPARATIVE TRUNKING REQUIREMENTS

(a) Full-Load Networks

Network	Trunk Miles		Trunks		$\bar{B}$ (Engineered)
	Actual (000)	% Diff. from 2-Level Hier.	Actual	% Diff. from 2-Level Hier.	
2-level hier.	1174	0	6659	0	0.007
3-level hier.	1154	−1.7	6679	+0.3	0.008
Symmetrical	1129	−3.8	6727	+1.0	0.007
Gateway	1268	+8.8	9236	+38.6	0.010

(b) Reduced-Load Networks

Network	Trunk Miles		Trunks		$\bar{B}$ (Engineered)
	Actual (00)	% Diff. from 2-Level Hier.	Actual	% Diff. from 2-Level Hier.	
2-level hier.	6047	0	3298	0	0.008
Symmetrical	5801	−4.1	3256	−1.3	0.008

per cent difference in trunk miles and 1.3 per cent difference in trunks between the symmetrical and hierarchical reduced-load networks.

The gateway network requires a much larger number of trunks and trunk miles than any of the others, reflecting the fact that many calls which in the other networks require only one link must use three in the gateway, and the fact that there is much excessive routing, or "back-haul" in traffic which is obliged to switch through gateways. In this case the resulting cost difference represents the savings in switching and line costs which would have to be achieved to offset the increased quantities of equipment required.

4.2 *Costs*

Table V gives the costs of the various networks, assuming a range of ratios of line to terminal costs which should include most actual situations. The differences between the hierarchical and symmetrical networks are quite small, as is that between the two- and three-level hierarchies, leading to a tentative conclusion that in these cases cost differential is not a primary reason for selection of one network over another. It must be remembered, however, that the hierarchy (two-level) was engineered using a known and proven economical procedure, while no such method is available for the symmetrical networks. Therefore, the hierarchies are probably close to optimal, while some additional economies might ultimately be realized for the symmetrical networks.

TABLE V — NETWORK COSTS

(a) Full-Load Networks

Cost/Trunk Mile	Using $1500 Trunk Termination and Switching Costs					
	$10/Trunk Mile		$50/Trunk Mile		$100/Trunk Mile	
Network	Cost $(000,000)	% Diff. from 2-Level Hier.	Cost $(000,000)	% Diff. from 2-Level Hier.	Cost $(000,000)	% Diff. from 2-Level Hier.
2-level hier.	21.72	—	68.69	—	127.39	—
3-level hier.	21.56	−0.74	67.72	−1.34	125.42	−1.55
Sym.	21.38	−1.57	66.54	−3.13	122.99	−3.45
Gateway	26.53	+17.54	77.25	+12.46	140.65	+10.41

(b) Reduced-Load Networks

2-level hier.	10.99	—	35.18	—	65.42	—
Sym.	10.69	−2.73	33.89	−3.67	62.89	−3.86

Using the $50 per trunk-mile line cost figure, the data shown in Table V(a) indicate that there is about a 1.3 per cent savings in cost of transmission and switching facilities for a network of this size with full loads when a three-level rather than a two-level hierarchy is used, and another 1.8 per cent if a symmetrical network is considered. This, of course, does not include any differences in signaling and control equipment which might be required to implement a symmetrical network, nor can it take account of the nonoptimality of the network engineering procedure now in use.

The gateway network, as expected, costs about 12 per cent more than a hierarchy to carry the same traffic, assuming that trunk and terminal costs are the same for all networks. This difference will be somewhat mitigated by the fact that trunk and facility routes are more likely to be identical in the gateway than in the hierarchical configuration, and therefore the multiplier to convert from airline miles to facility miles may well be smaller. In addition, any savings in switching costs which can be effected because of the large volumes of traffic flowing through the gateway switches will of course work to the advantage of the gateway plan.

Table V(b) indicates that the reduced-load symmetrical network is about 3.7 per cent less expensive than the hierarchy. This is in agreement with earlier results[1] which indicated that lightly loaded networks show greater differences between configurations than do heavily loaded networks.

4.3 *Overload Performance*

Fig. 8 shows the over-all average blocking probability, $\bar{B}$, for four different full-load network configurations (two-level hierarchy, three-level hierarchy, symmetrical, and gateway), and for the load changes discussed in Section 3.5 above. The "base" load under the "BH Runs" heading represents the same average load as the "engineered" point. It is a shorter run, however, and any difference in blocking between the two points is due to statistical fluctuations. Only the points on the charts are meaningful, but lines have been drawn connecting them for visual clarity. Fig. 9 is a similar chart, showing the dispersion factor, D. Figs. 10 and 11 show the same factors for the reduced-load networks, where only the two-level hierarchy and symmetrical networks were examined.

Although there are apparently some differences in performance between the various networks under various load change conditions, it is clear from Figs. 8 and 9 that there is no single superior network configuration in terms of traffic capacity and performance under shifting loads

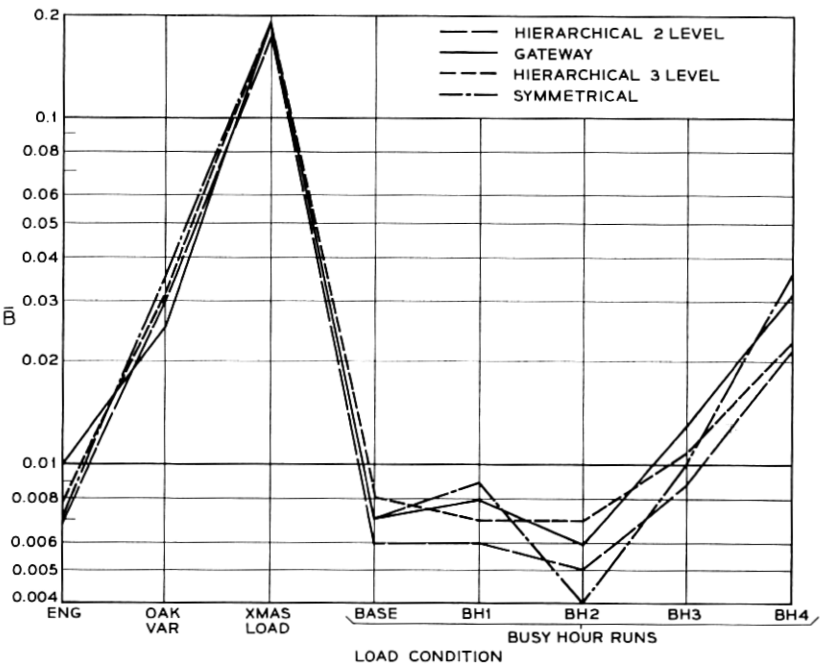

Fig. 8 — Over-all average blocking, $\bar{B}$: comparison of network configurations, full-load networks.

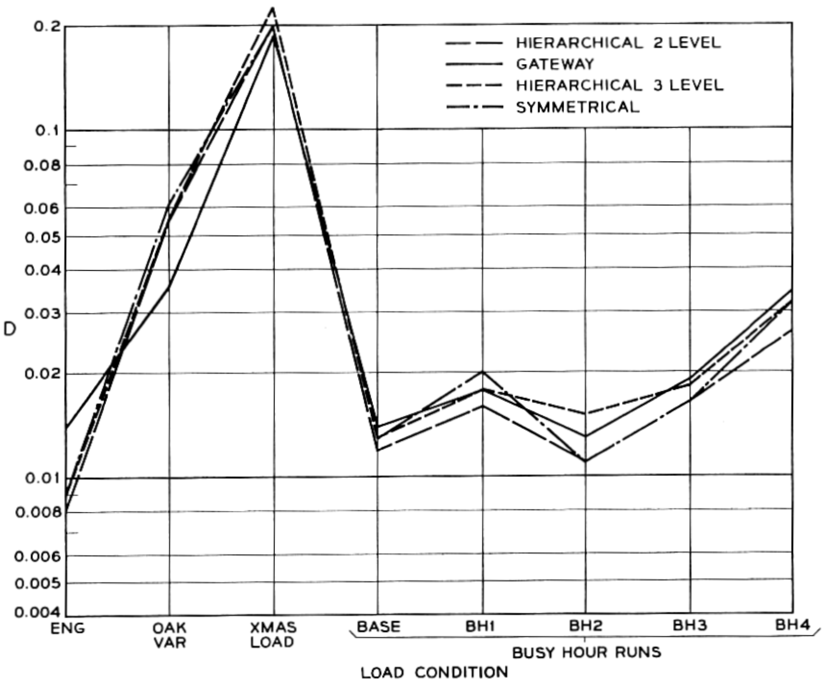

Fig. 9 — Dispersion of blocking, D: comparison of network configurations, full-load networks.

at full-load levels. Some small systematic differences are present, such as the fact that the two-level hierarchy appears to give slightly lower blocking than all other networks at all points except the Oakland variations case, where the gateway shows up best. This can, however, be a result of the initial engineered blocking level, which is slightly lower for the two-level hierarchy than for the three-level hierarchy or the gateway. This initial point does not so much denote a difference in performance under changed loads as it does the slight inaccuracies in engineering level, which are then reflected at every point on the chart. Although the simulation runs which produced these measurements used the identical set of calls for all networks at each load, the standard deviation of the results due to the finiteness of the simulation run is of the order of magnitude of the blocking probability at each point, and firm conclusions can be drawn only if a distinct superiority of one configuration over another manifests itself at almost all of the points considered. There are some such uniform results, but the differences are quite small, and may be offset by the differences in cost discussed above.

Figs. 10 and 11, on the other hand, show a small advantage for reduced-load hierarchical networks under all changed load conditions. In this case there is no initial error, and all evidence indicates that the hierarchy is slightly superior. It must be remembered, however, that the hierarchy costs somewhat more in this case, and this sensitivity to overloads may simply be the penalty paid for a more economical network at engineered loads.

The conclusion which must be reached from these results is that, for large networks with fairly high traffic densities, the performance of various alternate routing configurations in terms of traffic capacity under changing load conditions is quite similar. The reason for this is probably that the very density of traffic in these networks causes many of the trunk groups to be quite efficient, and the great bulk of the traffic is carried on the direct routes. Differences in alternate routing configuration, therefore, affect only a small proportion of the total traffic, with a correspondingly small effect on the network performance. In more lightly loaded networks, as has been observed, the differences are greater as

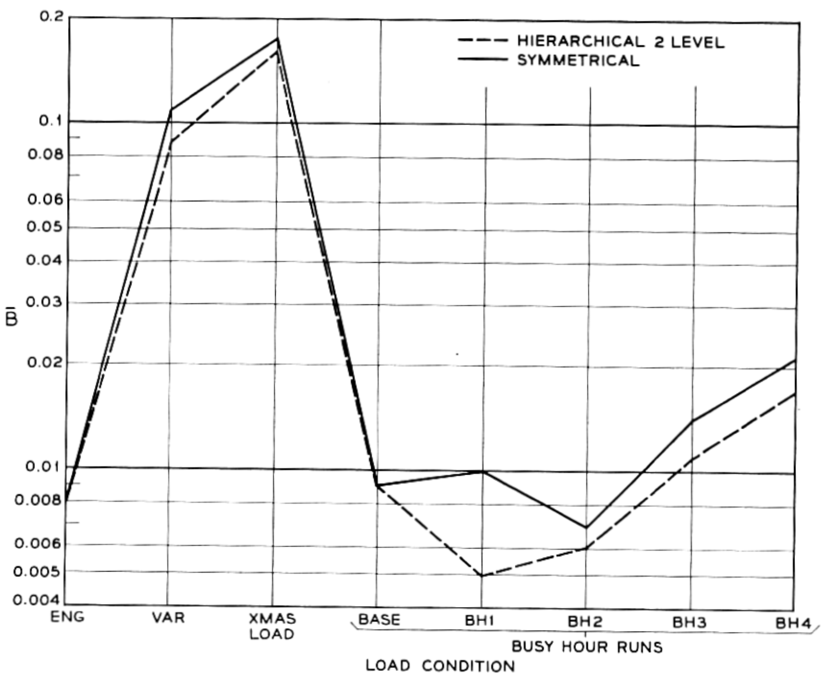

Fig. 10 — Over-all average blocking, $\bar{B}$: comparison of network configurations, reduced-load networks.

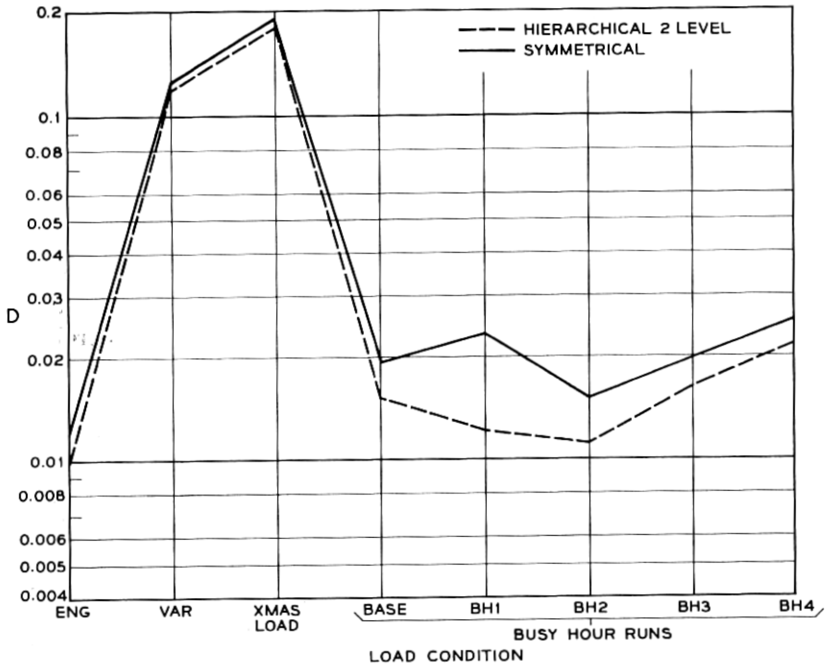

Fig. 11 — Dispersion of blocking, D: comparison of network configurations, reduced-load networks.

more of the traffic is alternate routed. Even in these cases, however, the differences are not large, and the comparison made here between symmetrical and hierarchical networks shows the slight superiority of one in cost to be offset by better performance of the other under shifting loads.

V. ANALYSIS OF OPERATING AND CONTROL PROCEDURES

5.1 *Full versus Limited Routing*

As discussed earlier, symmetrical networks were operated in two ways. In the first of these, called "limited routing," a maximum of three or four links per call was allowed, depending upon the connectivity available to the traffic parcel. In the second, called "full routing," five links per call were allowed for all calls. Fig. 12 shows a comparison of the over-all average blocking for these two cases. It is clear from this figure that operation with limited routing is superior in traffic handling capacity. Although the differences at any point are still small, and the statistical variability of the results large, the fact that there is an advantage for the limited routing case for every point tested indicates that this is a real effect, and not merely a result of chance observation. Furthermore, since these two curves represent the same network in terms of trunk layout, there is no possibility of complicating or compensating factors due to cost differences or engineering errors. In fact, the difference in blocking probability under engineered loads in this case does not represent an engineering error, but instead an additional verification of the fact that operation with limited routing is superior. This result is further evidence of the fact that excessive alternate routing can cause service deterioration, even under light load conditions. (The routing used in the symmetrical networks discussed earlier was limited routing, chosen because it gave superior performance.)

The symmetrical network whose performance is plotted in Fig. 12 and in subsequent graphs is clearly not identical to that discussed previously, since the blocking probabilities at all points are somewhat higher.

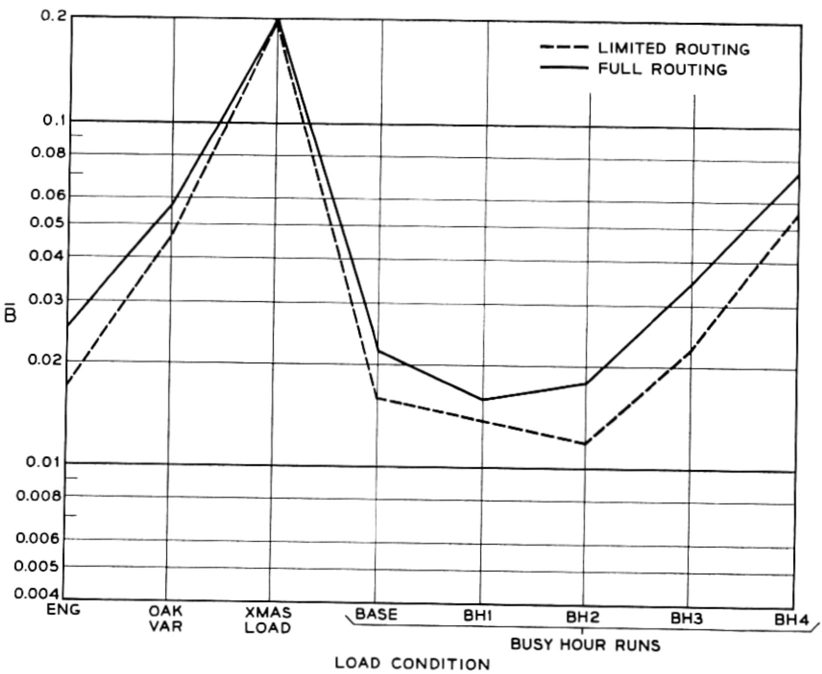

Fig. 12 — Over-all average blocking, $\bar{B}$: full vs limited routing; full-load, high-blocking, symmetrical network.

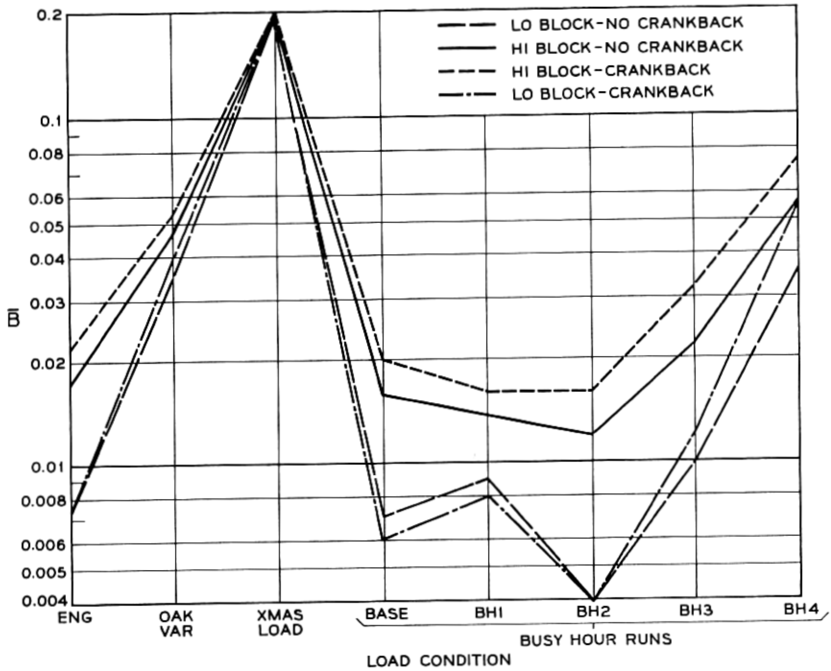

Fig. 13 — Over-all average blocking, $\bar{B}$: full-load symmetrical networks, effect of crankback.

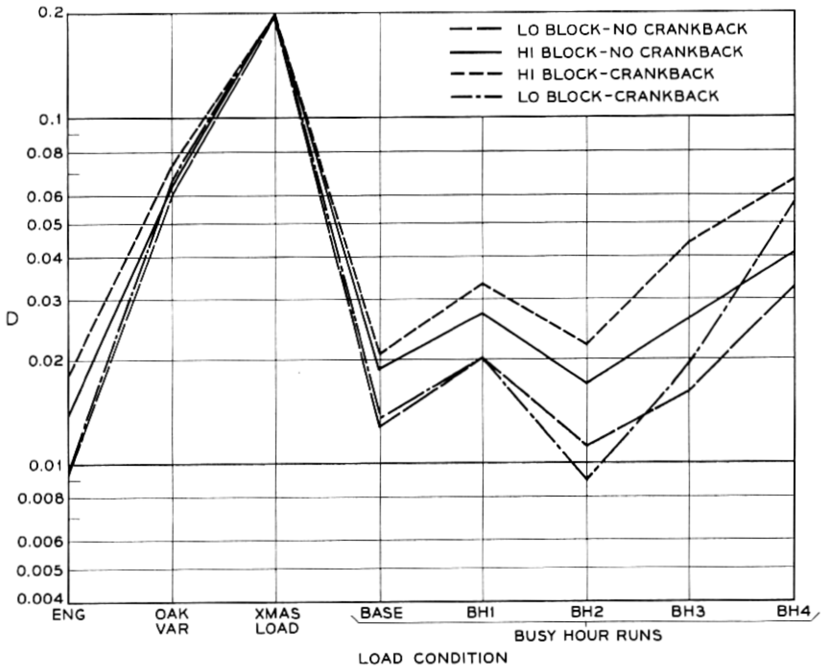

Fig. 14 — Dispersion of blocking, D: full-load symmetrical networks, effect of crankback.

This network, however, is the one which originally resulted from the engineering program, and it will be used for all studies concerned with differences in operating method using the same network. The earlier comparisons between network configurations on a basis of both cost and performance required that the blocking probability be approximately equal at engineered loads, and trial and error modifications were made to the symmetrical network to bring its blocking probability down to the proper level. The comparisons between different modes of operation of the same network should not be significantly affected by exact level of blocking at engineered loads and are expected to be valid for all networks of approximately the traffic densities considered. In the investigation of the effectiveness of crankback, however, symmetrical networks with both low and high over-all blocking probabilities were examined.

5.2 *Crankback*

Comparisons of networks operating with and without crankback were made for hierarchical and symmetrical networks using both the full traffic and reduced traffic models. In the hierarchical networks, no significant difference in behavior could be detected between the networks operated with and without crankback. This is because the structure of the hierarchical network is such that most of the blocking occurs on final route links, which are impossible to avoid even with the crankback option. For example, in Fig. 3(a) if a call from 1 to 2 is blocked at node 4, it may, with crankback, back up to node 1 and attempt to reach node 2. Even if this is possible, however, there is still a large probability of being blocked on route 3-2, and hence rearriving at 4 at some later time. In addition, those calls which do get through using the crankback option tend to use relatively long routes, causing later calls between other points to be blocked.

The over-all network blocking and dispersion of blocking for symmetrical networks with and without crankback are shown in Figs. 13 and 14 for full-load networks and in Figs. 15 and 16 for reduced-load networks. Figs. 13 and 14 have curves for both the symmetrical network as originally engineered (hi-block) and the corrected symmetrical network which was used for comparison with the hierarchy (lo-block). This was done so that any differences introduced by the general level of blocking would be apparent. The curves indicate that the level of blocking has, at most, marginal significance at these loads, and that crankback degrades the network performance at all but the lowest blocking levels, when it has virtually no effect. The networks here are operated with limited routing, but a similar test with full routing yields results which are substantially identical to those shown.

Figs. 15 and 16 show that crankback does offer a small advantage for less heavily loaded networks, although this advantage tends to disappear as the load increases, regardless of its distribution.

These results indicate that for large networks, operation with crankback at best offers a slight improvement in service when the service is good, and makes matters worse when the situation begins to deteriorate. An examination of the trunk occupancies and number of links per call shows that operation with crankback generally causes a larger number of links per call to be used on the average, with a higher over-all trunk occupancy. In effect, it therefore increases the amount of alternate routing allowed, and not always in the best way, so that degradation under overloads is a certainty. It therefore must be recommended that this device not be incorporated into large switching networks unless survivability, improved reliability, or other factors dictate it. If it is incorporated into a network for reliability or other purposes, means should be made available to disable it under overloads.

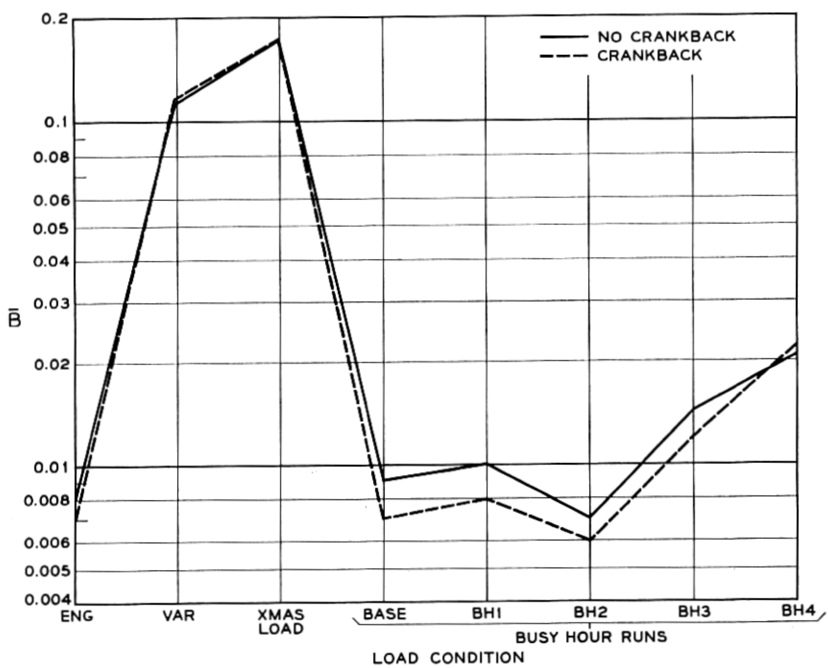

Fig. 15 — Over-all average blocking, $\bar{B}$: reduced-load symmetrical network, effect of crankback.

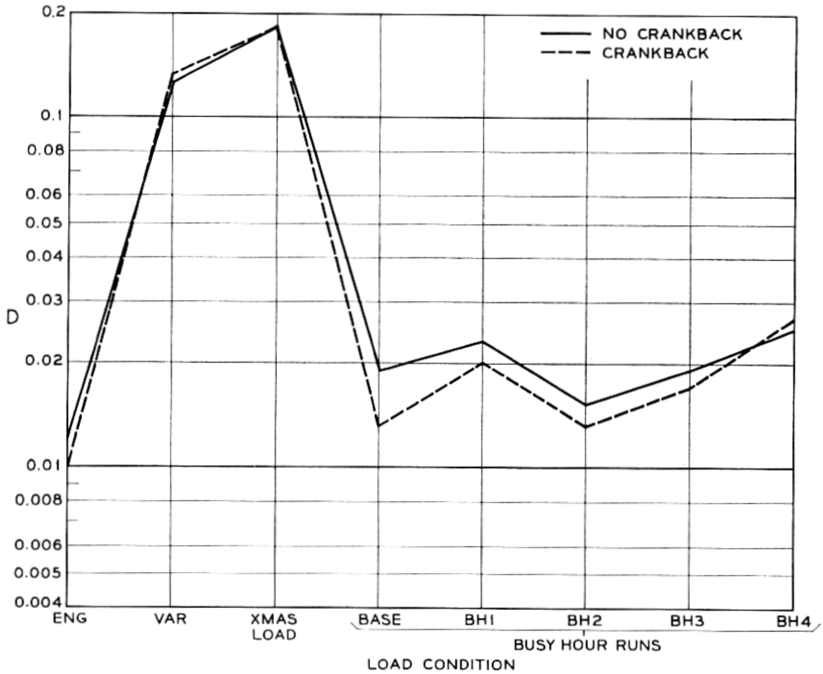

Fig. 16 — Dispersion of blocking, D: reduced-load symmetrical network, effect of crankback.

5.3 *Trunk Reservation for First-Routed Traffic*

Figs. 17 through 22 show the effect of trunk reservation for first-routed traffic on the blocking and dispersion of full- and reduced-load symmetrical and hierarchical networks. This measure, which reduces the amount of alternate routing on a selective basis, provides a uniform improvement in performance for all networks shown, although the improvement is more marked in the case of full-load than in reduced-load networks. The two-level hierarchies were not noticeably affected by the introduction of this measure.

In general, one trunk was reserved in each trunk group in the network, although two trunks were reserved on every group in some cases. It was generally found that reserving more trunks than noted in the charts had little additional effect upon the network performance. Figs. 17 and 18 show the effect of trunk reservation on symmetrical full traffic networks. It is interesting to note that the network with full routing has almost identical performance to the network with limited routing when trunk reservation is used. This is not illogical,

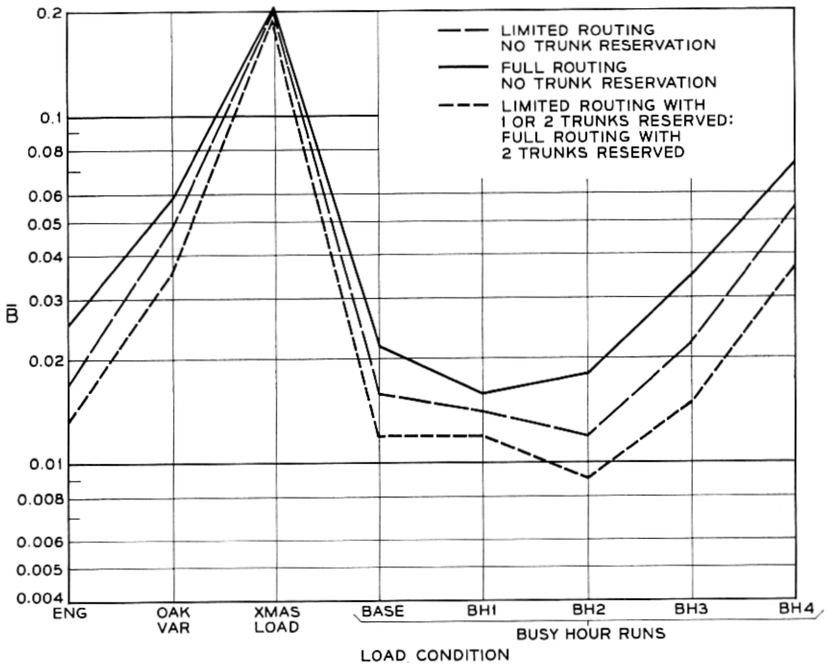

Fig. 17 — Over-all average blocking, $\bar{B}$: full-load, high-blocking symmetrical network, effect of trunk reservation.

since trunk reservation has a gross effect similar to that introduced by limiting the number of links per call.

Figs. 19 and 20 show the effect of trunk reservation on a three-level hierarchical network, and here we observe an improvement similar to that seen in the examination of symmetrical networks.

Figs. 21 and 22 show the blocking and dispersion for the reduced-traffic symmetrical network, in which the effect is similar but of lesser magnitude than that observed in the full-load networks.

It is quite likely that a selective application of trunk reservation to those groups which are large and have a large proportion of alternate routed traffic would be more effective than the across the board application used here. However, this study suffices to show that there is an advantage in the traffic handling capability of a network so equipped, and more detailed analysis will be required to determine the best number of trunks to be reserved in any given case.

Trunk reservation has essentially the opposite effect on the network as crankback; it reduces the amount of alternate routing during periods of momentary congestion, preventing calls from being completed using circuitous routes at such times. Subsequent calls are then not affected and the over-all network performance is improved.

One test was made using both trunk reservation and crankback, but the effect of trunk reservation appeared to dominate, and no difference was observed whether crankback was or was not used.

VI. CONCLUSIONS

The first and most obvious conclusion to be drawn from the preceding results is that for networks with a high traffic density the selection of routing doctrine and control philosophy does not have any great effect upon the traffic handling capability of the trunking network. This fact is apparently due to the substantial trunk group size generally encountered in such networks, with the basic group efficiency sufficiently large to obviate any spectacular improvements due to clever routing or control schemes. Of course, these comments apply only to reasonable alterna-

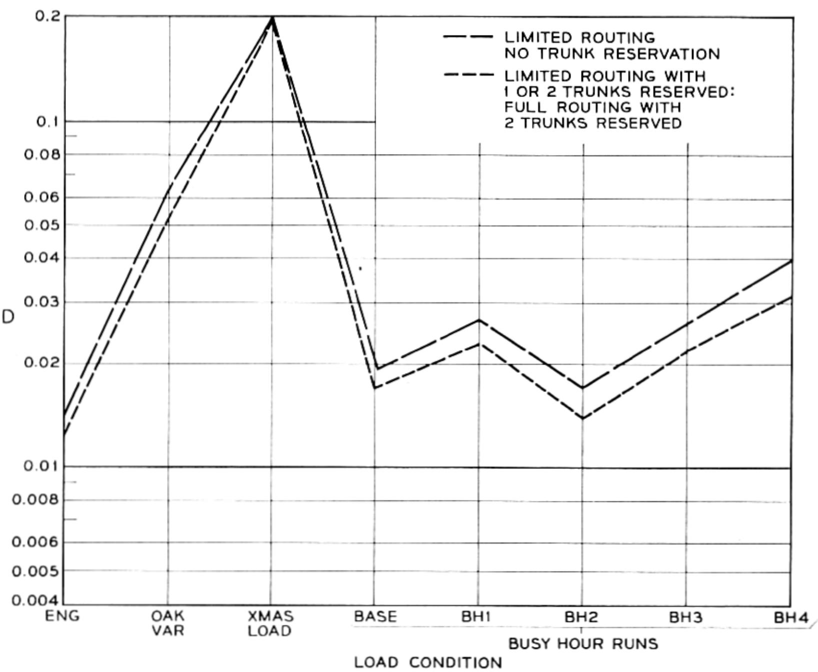

Fig. 18 — Dispersion of blocking, D: full-load, high-blocking symmetrical network, effect of trunk reservation.

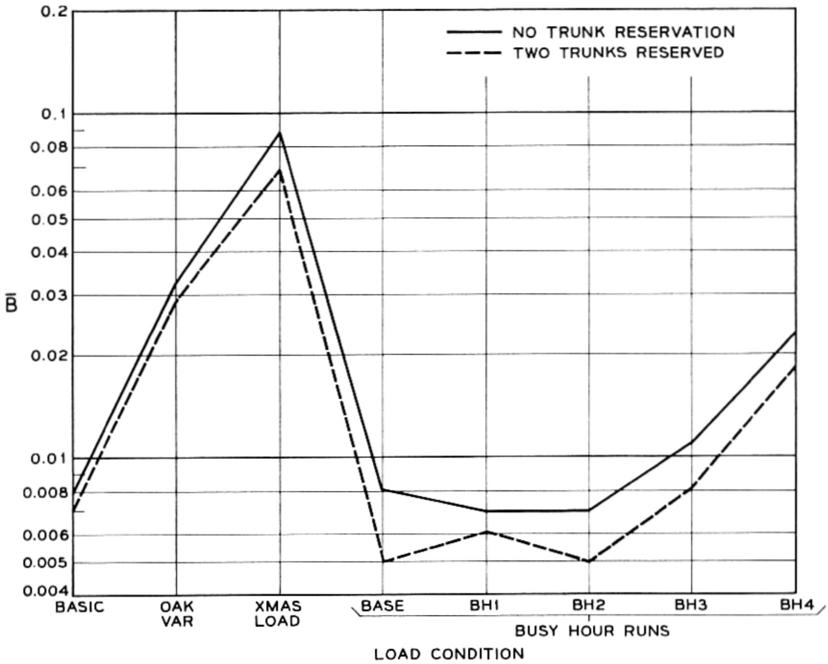

Fig. 19 — Over-all average blocking, $\bar{B}$: three-level hierarchical full-load network, effect of trunk reservation.

tives, such as those examined here. It is possible to develop a routing plan which would encourage circuitous routing at the expense of direct. Such a scheme would almost certainly show significantly poorer behavior than any of the networks investigated.

Planning for future networks should then be initially concerned with other factors, such as economics, survivability, flexibility and so forth, with a precise evaluation of traffic capacity to be determined after the fundamental design considerations are well formulated.

Having once accepted the basic idea that all differences are small in magnitude, we can nevertheless observe their direction, and, in the event that there are no other significant factors, decisions can be made on the basis of such small differences. A saving of one per cent in the toll trunk plant in the U.S.A. alone, for example, would amount to many millions of dollars, which is not insignificant in magnitude, even though it is a small fraction of the total network cost.

In the comparison of network configurations, the symmetrical networks have some cost advantages, particularly at lower load levels. This is to some extent offset by a tendency to deteriorate under overload slightly more rapidly than hierarchical or gateway networks. Furthermore, there is likely to be a not insignificant additional cost connected with the operation and control of such networks, and they are difficult to engineer and administer. They do have the advantage of improved survivability, however, since there is not so much concentration of facilities at regional switching centers.

The gateway network behaves well under overloads, but requires too high an initial cost to warrant its use with existing technology. If technological advances radically change the patterns of costs for such a network, then the gateway may be a suitable selection. The survivability aspects of these networks are particularly important, since sections of the network can be isolated by the destruction of a few critical points.

The hierarchical networks, which were the first alternate routing networks to be put into service, show a competitive initial cost and a reasonable reaction to shifting loads of all sorts. They are simple to engineer

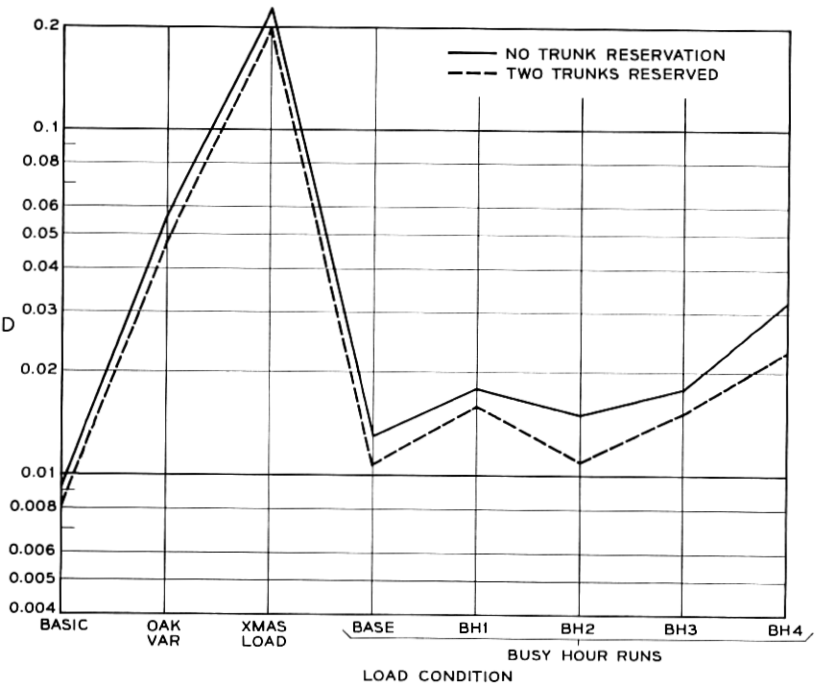

Fig. 20 — Dispersion of blocking, D: three-level hierarchical full-load network, effect of trunk reservation.

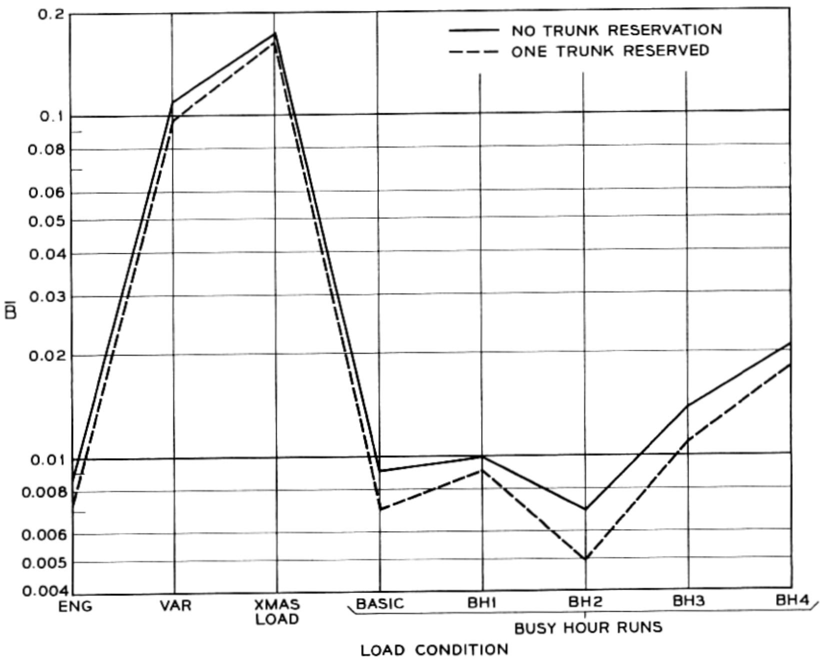

Fig. 21 — Over-all average blocking, $\bar{B}$: reduced-load symmetrical network, effect of trunk reservation.

and administer, and the logic associated with switching and routing control is relatively uncomplicated and economical. They pose an obvious survivability problem, since some traffic parcels have access to only a single route. This situation can be largely alleviated by dispersion of routes and liberal provision of high-usage groups.

In short, if a high-density communications network is desired, and concentration of traffic along backbone routes is allowable, then a hierarchical network is likely to be the best choice of network structure. As the traffic density declines, the symmetrical networks begin to show to advantage, and they are indispensable in some form if the survivability requirement rules out hierarchies. Symmetrical networks should, however, be implemented only in conjunction with an operating technique such as trunk reservation to maintain overload capability.

The investigations of control measures demonstrate conclusively that crankback is ineffective or harmful in all networks except perhaps those with extremely light traffic densities. It offers at most a small gain at engineered loads, and aggravates undesirable overload effects. There would therefore appear to be no reason for providing it other than the nontraffic one of improving the ability of a call to avoid an equipment malfunction. If it is used for this purpose it should be disabled under overload, when it shows the greatest traffic disadvantage.

Trunk reservation, on the other hand, almost always improves the traffic carrying capacity of networks, and is never harmful. It is an inexpensive measure to implement which is unquestionably worth using, and further studies of the strategy and extent of its use should be undertaken.

In sum, the basic factors relevant to the design of communications networks are:

(1) If there is a high density of traffic, and traffic concentration on backbone routes is allowed, then a hierarchical configuration probably should be selected, with the number of levels dependent upon the particular situation.

(2) If the traffic density is lower and/or the hierarchy is unacceptable

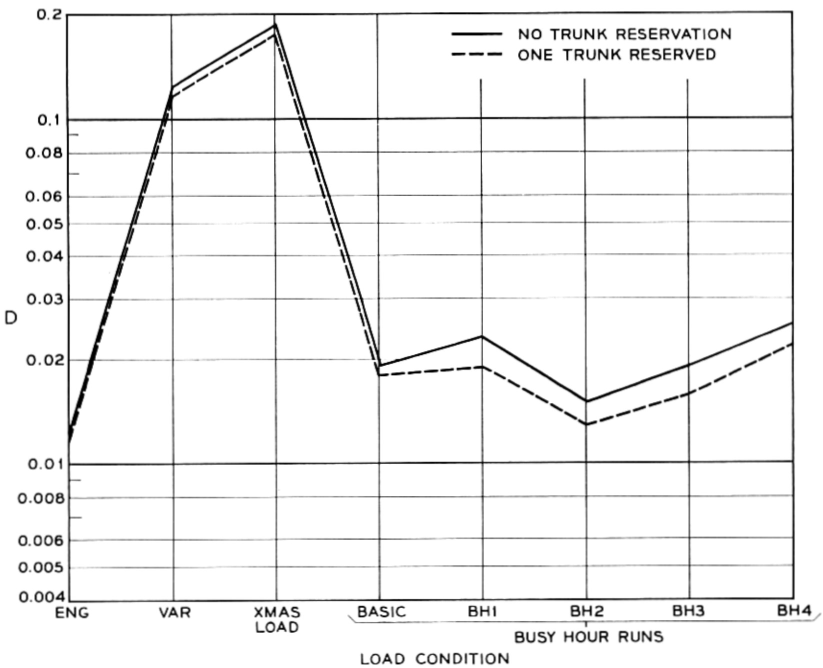

Fig. 22 — Dispersion of blocking, D: reduced-load symmetrical network, effect of trunk reservation.

for survivability reasons, then a symmetrical network may be more economical and can perform well if properly controlled.

(3) Crankback should not be used, except possibly as a means of alleviating the effects of equipment troubles. If used, its traffic disadvantages under overloads should be taken into account.

(4) Trunk reservation should be widely employed, since it is simple to implement and has noticeable traffic advantages under all load conditions with almost any network configuration.

Although these guidelines are, of course, qualitative in nature, this is necessary because of the large number of variables which exist in an actual network. Variations in traffic levels between and within networks, geographical distributions of switching offices and densities of traffic, equipment limitations and differing primary functions all lead to different constraints and weightings of various factors. It is the purpose of these studies to provide guides for the design of communications networks, with final choices dependent upon specific factors.

VII. ACKNOWLEDGMENTS

This work could not have been done without the efforts of L. A. Gimpelson in organizing, and Mrs. E. E. Bailey and Misses G. C. Watling, S. A. Switch and A. Malone in programming, the simulation. Miss M. Lynch prepared most of the voluminous input data.

REFERENCES

1. Weber, J. H., Some Traffic Characteristics of Communications Networks with Automatic Alternate Routing, B.S.T.J., **41**, Mar., 1962, pp. 769–796.
2. Gimpelson, L. A., and Weber, J. H., UNISIM—A Simulation Program for Communications Networks, Fall Joint Computer Conference, San Francisco, Oct., 1964.
3. Pilliod, J. J., Fundamental Plans for Toll Telephone Plant, AIEE Trans., **71**, pt. I, 1952, p. 257; also in B.S.T.J., **31**, July, 1952, p. 832.
4. Truitt, C. J., Traffic Engineering Techniques for Determining Trunk Requirements in Alternate Routing Trunk Networks, B.S.T.J., **33**, Mar., 1954, p. 277.
5. Wilkinson, R. I., Theories for Toll Traffic Engineering in the U.S.A., B.S.T.J., **35**, Mar., 1956, pp. 421–514.
6. Segal, M., Traffic Engineering of Communications Networks with a General Class of Routing Schemes, Fourth International Teletraffic Conference, London, July, 1964.