THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 54, No. 5, May–June 1975
Printed in U.S.A.

Digital Data System:

Network Planning

By P. F. BROWN, JR., G. W. PHIPPS, and L. A. SPINDEL

(Manuscript received July 12, 1974)

Network design methods are described which determine the multiplexing and digital transmission facilities required to serve a given data-circuit demand. The long-haul network design is based on the definition of a three-level network routing hierarchy, the derivation of intercity bit-stream requirements, and a technique for selecting digital transmission facilities to carry the bit streams. The distribution network design within metropolitan areas makes use of different multiplexing combinations best suited to anticipated demands and techniques for using short-haul digital facilities in an efficient manner.

I. INTRODUCTION

Given the basic two-stage multiplexing scheme to be used in the DDS and a forecast of data-circuit demands at a set of nodes, or cities, the basic purpose of DDS network planning is to define the network in terms of the multiplexing equipment required at each node and the arcs, or digital transmission facilities, required to interconnect the nodes. This network definition can then be used directly to estimate the manufacturing requirements for new equipment, to determine locations between which installation of digital transmission facilities is required, and to calculate the capital resources needed to implement the network.

This paper describes the methods used to define the DDS network. Sections II, III, and IV focus on the long-haul, or intercity, part of the network, its structure and hierarchy, the multiplexing algorithm used to determine nodal multiplexer arrangements and internodal bit-stream cross sections, and the algorithm used to route digital bit streams over existing or planned digital transmission facilities. Section V describes the network arrangements within cities, or the local distribution part of the DDS network.

The major goal of the DDS long-haul network design is to maximize transmission efficiency, or digital bit-stream fills, while minimizing multiplexer cost. This implies collecting and routing individual customer circuits, ranging in speed from 2.4 to 56 kb/s, in such a way as to economically trade off multiplexer cost with efficient use of each DS-1 bit stream. Also implied is the need to efficiently utilize the available low-cost long-haul bit-stream facilities, namely the DS-1 channels derived from the application of the 1A Radio Digital System to existing radio routes. The algorithm described in the following sections has proven quite successful in meeting this goal. With typical data market forecasts, transmission fills are expected to be 70 percent of capacity by the end of the second year of service.

The primary purpose of the local distribution network design is to minimize multiplexer costs. This is accomplished by clustering the data circuit demands that can be served by a single multiplexer and tailoring each multiplexer arrangement to meet the demand projected for it.

II. NETWORK HIERARCHY

A three-level network hierarchy has been defined for the DDS. The first step in defining the hierarchy is to select those cities, or digital serving areas (DSAs), that are the highest level in the hierarchy; these are called Class I DSAs. Given a list of about 100 cities that are to be served on the network during the first three to four years of service, the Class I DSAs were selected based on the following criteria:

(*i*) Data-circuit demand. Class I DSAs are those cities which have a relatively large number of data circuits that could potentially be served by the DDS.

(*ii*) Geography. At least one city in each major region of the country is designated as a Class I DSA.

(*iii*) Transmission facility access. Class I DSAs generally are those cities that have access to large cross-section transmission facilities and are thus capable of "collecting" data circuit demands in a large region for transmission to other regions.

Once the Class I DSAs were selected, Class II DSAs, each of which homes on a single Class I DSA, were designated based on data-circuit community of interest and geographic proximity. Finally, Class III DSAs, each of which homes on a single Class II DSA, were designated based primarily on their location relative to higher-level DSAs in the network.

Class I DSAs are connected to other nearby Class I DSAs in such a way as to take advantage of existing major transmission-facility cross sections. These arcs between Class I DSAs, together with the homing arcs connecting Class II DSAs to Class I DSAs and Class III DSAs to

Class II DSAs, form a fully connected network. These arcs make up the backbone network through which a connection for any circuit can be guaranteed. After many trials with a network simulator were made, the network hierarchy and backbone connections shown in Fig. 1 for an example 96-city network were selected.

Not all digital bit streams in the DDS are routed on the backbone network. If the data-circuit demand between any DSA pair is such that a digital bit stream can be filled to a given level, then that bit stream will not be demultiplexed at intermediate DSAs. The fill parameter currently being used in determining these express bit streams is 70 to 80 percent; that is, if the total number of data bits to be transmitted per unit time exceeds 70 to 80 percent of the data-bit capacity of the digital bit stream, then an express bit-stream requirement is defined. The express bit-stream concept is applied both to DS-0 bit streams between submultiplexer pairs and to DS-1 bit streams between T1DM pairs to efficiently utilize both multiplexers and facilities.

III. BIT-STREAM MULTIPLEXING

Bit-stream multiplexing is the planning function that transforms the 2.4-, 4.8-, 9.6-, and 56-kb/s data-circuit demands between DSA pairs into DS-0 and DS-1 bit-stream requirements between DSA pairs, subject to the hierarchy definition and express bit-stream concepts described above. The DS-0 requirements determine the number and location of submultiplexers required at nodes in the network. The DS-1 bit-stream requirements determine the number and location of T1DMs and are also used as input to the routing process described in Section IV. A basic engineering rule is that multiplexers can be located only at the test center locations serving a given city; this precludes locating multiplexers at intermediate points such as radio junctions. Thus, the only nodes considered in the bit-stream multiplexing portion of long-haul network planning are the test center locations in each DSA.

The inputs to the multiplexing algorithm are four matrices representing the 2.4-, 4.8-, 9.6-, and 56-kb/s data circuit demands between nodes and the hierarchical definition of the nodes. An element $A_{i,j}$ in a given matrix thus represents the number of data circuits forecast between nodes i and j.

To demonstrate the multiplexing algorithm used, consider the simplified example in Fig. 2. Assume that nodes 1 through 4 are Class III and nodes 5 and 6 are Class II. Further, assume that the demand for 2.4-kb/s circuits between node pairs is given in Fig. 3a.

The problem then is to transform the matrix in Fig. 3a into a similar matrix giving the number of DS-0 bit streams required to carry the

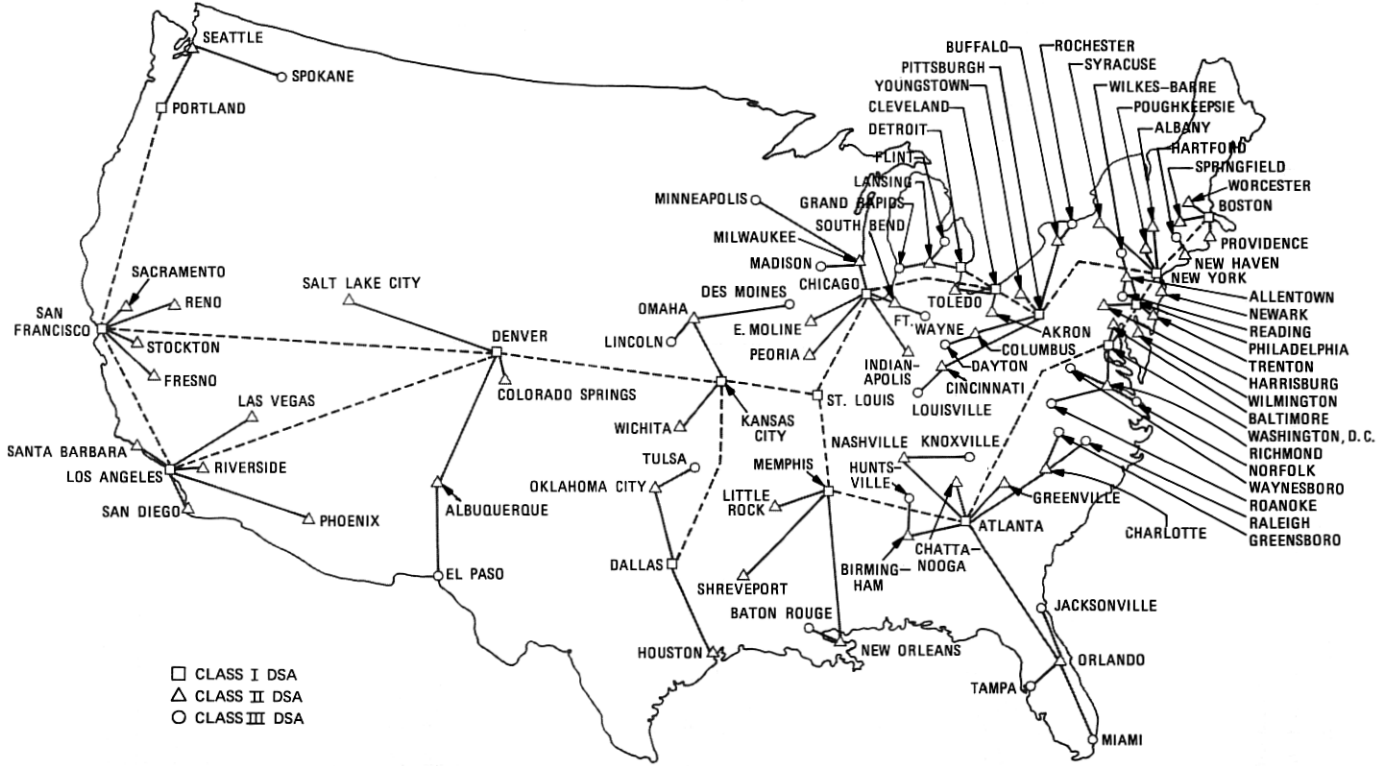

Fig. 1—Projected 96-city network.

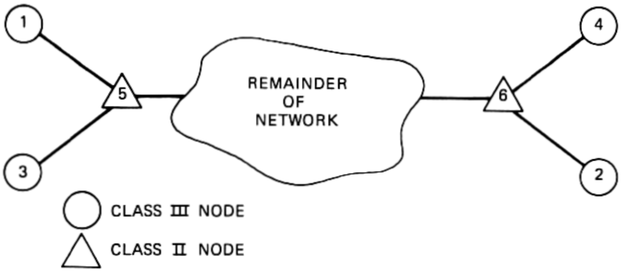

Fig. 2—Multiplexing algorithm example.

2.4-kb/s circuits between node pairs. Two parameters are important in the transformations. The first is the capacity C_r of the DS-0 bit stream, where r is the circuit speed. Based on the submultiplexing plan, $C_{2.4}$ is fixed at 20 for 2.4-kb/s circuits. The second parameter specifies the minimum fill F_r at which express DS-0 bit streams are established. For purposes of this example, assume that a minimum fill of 70 percent is required so that $F_r = 0.7$ and $F_rC_r =$ fourteen 2.4-kb/s circuits for an express 2.4-kb/s bit stream.

The processing of the demand matrix starts with the leftmost element of row 1 and proceeds element by element across the row.

	1	2	3	4	5	6
1		24	8	42	6	16
2			4	12	18	10
3				6	0	2
4					2	2
5						4
6						

(a)

	1	2	3	4	5	6
1		20	8	42	10	16
2			4	12	22	10
3				6	0	2
4					2	2
5						4
6						

(b)

	1	2	3	4	5	6
1		20	0	40	20	16
2			4	12	22	10
3				6	8	2
4					4	2
5						4
6						

(c)

	1	2	3	4	5	6
1		20	0	40	20	16
2			0	0	20	28
3				0	20	0
4					0	24
5						22
6						

(d)

	1	2	3	4	5	6
1		1	0	2	1	1
2			0	0	1	2
3				0	1	0
4					0	2
5						2
6						

(e)

Fig. 3—Transformation of 2.4-kb/s circuits to 64-kb/s bit streams.

Consider element $A_{i,j}$, where $i = 1, j = 2$. Find the largest n for which

$$A_{1,2} - nC_r \geqq 0$$

or

$$\begin{aligned} 24 - n20 &\geqq 0 \\ n &= 1. \end{aligned}$$

Now determine

$$A_{1,2} - nC_r = R.$$

If $R = 0$ or $R \geqq F_rC_r$, or if node i homes on node j, leave the $A_{i,j}$ unchanged in the original demand matrix. Otherwise, enter nC_r in the place of the $A_{i,j}$ being considered and add R to the $A_{i,k}$ and $A_{\min(j,k),\max(j,k)}$ elements, where k is the designator for the node on which node i homes in the hierarchy. In the case shown,

$$R = A_{1,2} - nC_r = 4.$$

Since $4 < 14$, 20 is entered in place of $A_{1,2}$ and 4 is added to $A_{1,5}$ and $A_{2,5}$. This operation gives the modified matrix shown in Fig. 3b. Repeating the same operation for the remaining four elements in the first row gives the matrix shown in Fig. 3c. The same element-by-element treatment is applied to each row in order, always beginning with the leftmost element in the row until the final matrix is the new representation of 2.4-kb/s demand, modified to account for the network hierarchy and express routing. The fully modified matrix for the example is shown in Fig. 3d (ignoring the remainder of the network beyond the nodes shown). This matrix is transformed into the DS-0 bit-stream matrix shown in Fig. 3e, in which the elements

$$\beta_{i,j} = \left[\frac{A_{i,j}}{C_r} \right].$$

The brackets denote the integer greater than or equal to

$$\frac{A_{i,j}}{C_r}.$$

Note in the above process that, if $R = 0$ or $R \geqq F_rC_r$, all the demand for the $A_{i,j}$ being considered is carried on express or high-usage bit streams between nodes i and j. If there is some $R < F_rC_r$, that portion of the demand is carried on a backbone network bit stream one level up the hierarchy and reconsidered for express routing at the higher-level node.

The process described is followed for all nodes in the network. However, when $R < F_rC_r$ for Class I nodes, there is no higher network level, and a look-up table is used to determine the adjacent node to which the partially filled bit stream is routed. For example, Atlanta demand may be routed to either Memphis or Washington, depending on its ultimate destination. Several iterations through the Class I nodes are necessary to determine DS-0 bit-stream requirements.

A matrix transformation process identical to that described above for 2.4-kb/s circuits is made for 4.8- and 9.6-kb/s demand matrices, with appropriate values for C_r and F_r. The 56-kb/s data-circuit demand matrix does not require transformation, since each 56-kb/s data circuit takes the capacity of one DS-0 bit stream. Note that the transformed matrices give the submultiplexers of each speed required at each node in the network.

The matrices giving the DS-0 bit-stream requirements for the four speeds are then added together to give the total requirements between each node pair. The same algorithm is then applied to this matrix to transform it into a matrix showing requirements for DS-1 bit streams. Values are assigned C_r and F_r to obtain efficient transmission fills, while leaving some capacity for growth. The number and location of T1DMs can be derived directly from the transformed DS-1 bit stream matrix.

Observe that the bit-stream multiplexing algorithm tends toward a network solution with relatively high transmission fill. Each express arc is, by definition, filled to at least 70 to 80 percent of its capacity. The only nonexpress arcs are on the backbone network and connect adjacent nodes. Since the backbone arcs can carry any combination of nonexpress circuits, regardless of orginating and terminating nodes, they are generally also used efficiently.

IV. BIT-STREAM ROUTING ON LONG-HAUL FACILITIES

The DS-1 bit-stream requirements, determined as described in Section III, are applied to the network of long-haul digital facilities. The primary vehicle for DS-1 bit-stream transmission in the early years of the DDS network is the 1A Radio Digital System (1ARDS), which applies one DS-1 bit stream to a microwave radio channel in combination with multiplexed voice circuits. The supply of these systems is limited to the number and location of radio channels used for voice transmission. It is then possible that the demand for DS-1 bit streams between two points in the network will exceed the radio system capacity for 1ARDS channels; in these cases, the bit streams are applied to larger capacity systems, such as an L-mastergroup digital system (LMDS).

4.1 Routing algorithm

The facility universe on which DS-1 bit streams may be routed includes about 300 nodes and 900 arcs. Some arcs will include facility capacity for more than one DS-1 bit stream. To route a DS-1 bit stream between two nodes, it is necessary to find some combination of arcs in tandem that has the capacity available to carry the bit stream.

A minimum-distance algorithm defined by Dijkstra[1] is used to find the minimum total facility "length" between two given nodes. It differs from others which find the minimum paths from one node to all other nodes. The technique is particularly attractive when paths are sought between pairs of nodes that are relatively close together.

Each node is assigned a two-dimensional quantity that indicates the homing node and the cumulative distance to the start (all distances are assumed to be positive). Beginning with the start node, the distances to nodes one arc away are calculated. These nodes then are assigned the node they home on and the cumulative distance to the start. Figure 4a is an example. Node A is the start node, and the shortest path to node E is sought. In the first step, nodes B and C can be reached from node A. In Fig. 4b, nodes B and C have been assigned to home on node A with distances to node A of 4 and 2, respectively.

The node with the shortest distance to node A is made "permanent." This means that the shortest path from the start to this node has been found. Since there are no negative distances, there are no possible shorter paths. In Fig. 4c, node C is made permanent. Distances to nodes one arc away from node C are computed. If the cumulative distance to the start for nodes reached on this new path is smaller than the previous value for that node or if the node has not been reached yet, then the new homing node and distance are assigned. Thus, node B is rehomed on node C, since the distance to the start is only $2 + 1 = 3$ through C, as opposed to 4 directly from A. Node F is $2 + 2 = 4$ away from node A and homes on node C. Node E is $2 + 8 = 10$ away from node A and homes on node C.

The node now with shortest cumulative distance to the start is node B (3 units from A as opposed to 4 and 10 for nodes F and E). Therefore, node B is made permanent, and distances to nonpermanent nodes one arc away are computed (see Fig. 4d). The cumulative distance to node E is $3 + 3 = 6$ by way of node B; this is smaller than the previous value (10) and thus represents a shorter route. Node E is homed on B with a distance of 6. Node D homes on B with a distance of 8.

The next node with shortest cumulative distance to the start is node F (4 units from A). Node F is made permanent (see Fig. 4e). Cumulative distances to nonpermanent nodes are computed. The distance to node E is $4 + 4 = 8$, which is *longer* than the previous

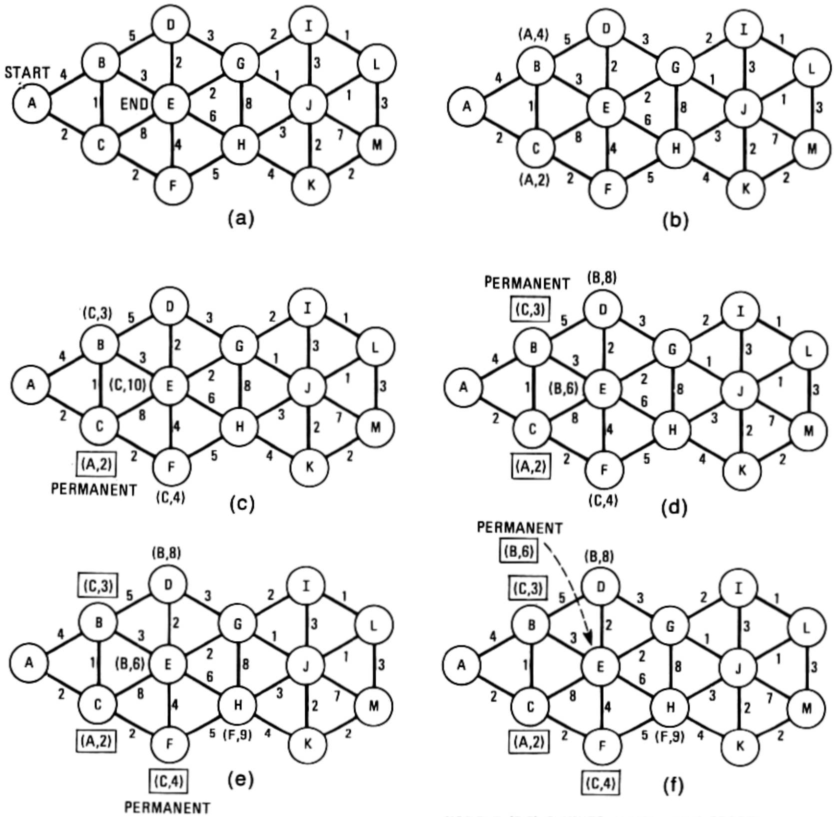

NODE E (B,6) 6 UNITS AWAY FROM START
NODE E HOMES ON B
NODE B HOMES ON C
NODE C HOMES ON A

Fig. 4—Dijkstra minimum-distance algorithm.

value to E, so E remains homed on B. Node H becomes homed on F with a distance of $4 + 5 = 9$.

Node E is the next nonpermanent node with the shortest distance to node A. Therefore, node E is made permanent, and the shortest route to node E from node A has been found (see Fig. 4f). The distance to the start is 6, and the path is found by tracing back through homing nodes: E to B to C to A.

In the DDS network, the primary engineering criterion for selection of a best path is to choose the route with the least number of arcs in tandem; a secondary criterion is to choose the shortest path in terms of route mileage. To apply these criteria within the Dijkstra algorithm, each facility is assigned a "length" equal to a relatively large constant plus a factor proportional to route mileage. The large constant ensures that primary consideration is given to the number of facilities required to route the bit stream, and the mileage factor gives secondary consideration to overall path length.

The algorithm is applied to each DS-1 bit stream individually in the order presented to the network model. As facilities are used, they are removed from the file and thus not considered for subsequent bit streams. This process continues until all bit streams are routed. If insufficient capacity is found at any point to route the bit stream being considered, this is noted. It should be clear that the order in which bit streams are routed affects the final network layout, as well as the particular facilities used to route a given bit stream. The flexibility allowed in the ordering of bit streams and the inherent ability to apply engineering judgment to specific routing problems is felt to be more desirable than a more rigid approach that considers the total set of bit-stream routing demands simultaneously.

4.2 Facility selection

The primary facility for inter-DSA transmission is the 1ARDS. However, as mentioned above, the 1ARDS will be applied only to existing radio systems, and sufficient capacity for the DS-1 bit-stream demand may not exist on some routes. A computer program is used to determine if a given data circuit demand can be routed entirely on 1ARDS channels and, if not, where higher capacity systems such as LMDS are required. The facility file is then augmented as required to reflect the need for the larger capacity systems.

For any arc in the network with capacity for more than one DS-1 bit stream, the facilities are selected in a predetermined order. Channels on high-capacity systems are selected first since their need has been previously demonstrated and they are generally more economic for large cross sections. 1ARDS facilities for a given arc are selected in an order to minimize voice-circuit rearrangement costs. Most radio systems will require no mastergroup modification before a 1ARDS signal is applied, and these facilities are selected first when 1ARDS channels are used on a given arc. Some radio systems require moving voice supergroups and modifying voice multiplexing equipment prior to application of the 1ARDS signal, and these facilities are selected only after other system capacity has been exhausted.

4.3 Example

Figure 5 shows the Class I node subset of an example 60-city network that contains a total of about 11,000 data circuit segments. The number

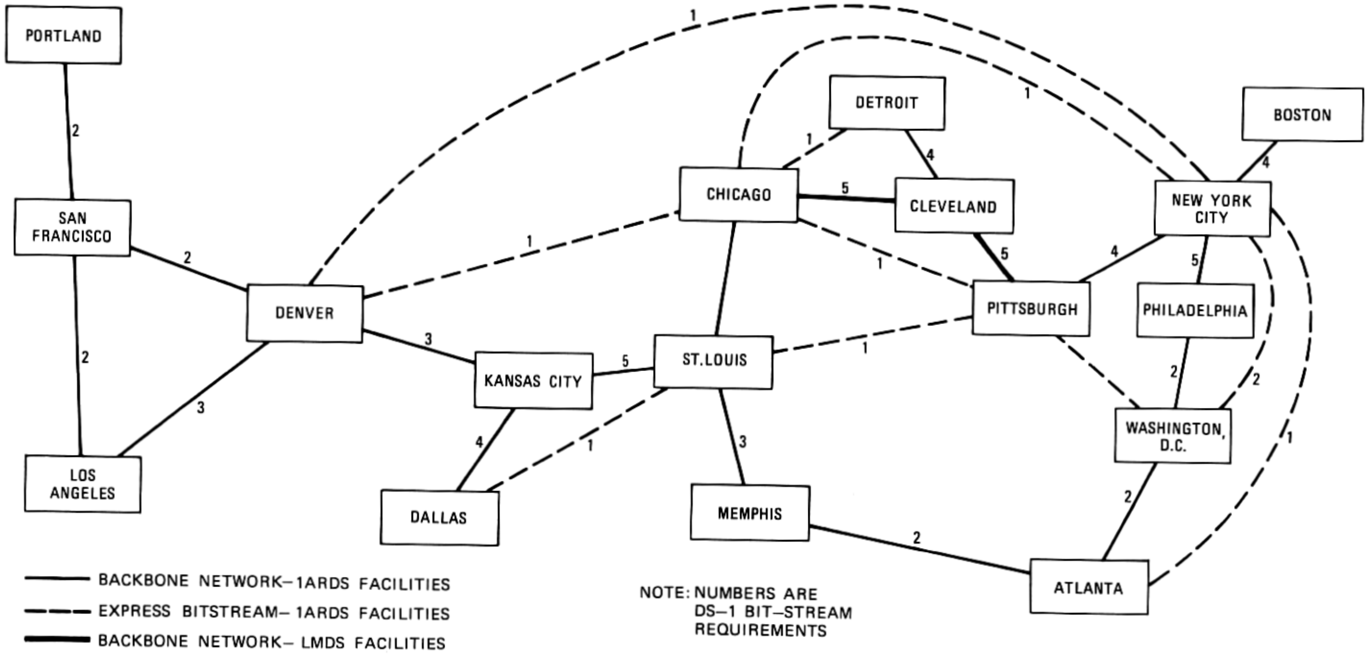

Fig. 5—Example network subset—backbone and express bit streams.

of DS-1 bit streams between Class I nodes is indicated for each arc in the network; a total of 155 DS-1 bit streams are required in the entire network. Note that in this case relatively few express DS-1 bit streams are generated in the network model.

This figure also shows that demand on only two network arcs exceeded radio capacity for the 1ARDS: Chicago to Cleveland and Cleveland to Pittsburgh. One LMDS channel was used on each arc to carry five DS-1 bit streams.

Figure 5 shows only the DS-1 bit-stream end points required in the network routing. It does not show the actual routing over existing and planned radio facilities, since sufficient detail could not be clearly shown. For example, the arc shown directly from Chicago to New York is actually routed from Chicago to Grant Park, Illinois, from Grant Park to Colesville, New Jersey, and from Colesville to New York. Also not shown is the bit-stream multiplexing information that specifies which DS-0 bit streams are carried on each DS-1 bit stream.

V. DIGITAL SERVING AREA DISTRIBUTION

Within a metropolitan area, which may include all central offices within about 50 miles of a test center, network planning takes on a different character. Here, the digital transmission distances are relatively short and so the trade-off between node costs and transmission costs is initially aimed at reducing node costs. This is done by attempting to minimize the number of multiplexers required and to tailor the multiplexer capacity to the anticipated demand. A secondary objective is to use T1 lines more efficiently through voice sharing or shared use by more than one data multiplexer.

The nodes in the DSA network are the set of central offices that are to be provided DDS service at a point in time. Since all data circuits originate in one of these central offices and must have a baseband appearance at the test center, the four data-circuit demand matrices have dimension $N \times 1$, where N is the number of central offices to be served; that is, the number of data circuit segments of each speed between each central office and the test center is the network demand. This is true for both the end sections of inter-DSA circuits, as well as intra-DSA circuits, since all must have the test center appearance.

The other major inputs required in the DSA planning process are the T1 carrier facilities available in the area, the length and gauge of cable pairs connecting central offices, and the length and gauge of loops served from the offices. In most cases, T1 and cable routes can be found between any two central offices, and this availability will be assumed in the following discussion.

5.1 An initial distribution network

A feasible solution to the DSA network design is to provide enough T1DMs and SRDMs at each node to serve the demand and to connect each node to the test center by a T1 line. (Recall that each working T1 line has a standby line; this does not reflect on the following process, and so only the working line will be referred to in the following description.) If the test center itself is coincident with one of the N nodes to be served, this feasible solution requires at least $N - 1$ multiplexer pairs and $N - 1$ T1 lines.

The first step in reducing node multiplexing costs is to find nodes whose demand can be served by a multiplexer at a nearby node. Consider a node j, for which

$$l_j + d_{ji} \leqq L,$$

where

l_j is the loop length at office j
d_{ji} is the interoffice cable length connecting office j to office i, and
L is the baseband transmission limit at the transmission rate being considered.

Then the demand at node j can be served by a multiplexer at node i and the multiplexer at j removed. The equation above must, of course, consider loop and cable gauges and is sensitive to transmission speed.

Using the above equation, all potential node clusters that can be served by a single multiplexer are enumerated. The cluster that serves the greatest demand is removed from the list, a multiplexer assumed at the corresponding node, and a second enumeration of clusters made with the nodes included in the first cluster removed. The second cluster is then removed based on the greatest demand served. In this way, a list of node clusters is formed in the order in which they will capture the greatest portion of data circuit demand. Each cluster can be served by a single multiplexer node, which must have the capacity to handle the combined data circuit demands of all nodes in the cluster.

In the DDS distribution network design, the transmission limit at 9.6 kb/s has generally been used to define limits on node clusters. This assures that circuits at the three lower data rates can be routed through their serving node to the node containing a multiplexer. For 56-kb/s circuits, a regenerator is being developed for use in cases where a station exceeds the 56-kb/s transmission distance from a multiplexer. This approach reduces the total number of multiplexers significantly and requires relatively few regenerators to reach outlying 56-kb/s stations.

A second step in reducing node cost is based on an estimate of future multiplexing capacity required at the node. As an example, suppose the demand on a multiplexer consists of 50 percent 2.4-kb/s, 20 percent 4.8-kb/s, 20 percent 9.6-kb/s, and 10 percent 56-kb/s circuits. For this speed mix, a T1DM/SRDM combination can multiplex about 120 circuits. With the same mix, a T1DM/ISMX combination can multiplex about 80 circuits at a lower per-circuit cost. Therefore, for nodes which have an estimated demand of less than 80 circuits over a reasonable time period, the provision of T1DM/ISMX multiplexing will reduce node cost with no increase in transmission cost.

A typical DSA network layout is shown in Fig. 6. Note that the number of nodes with multiplexers has been appreciably reduced by forming clusters. There is also widespread use of T1DM/ISMX multiplexing.

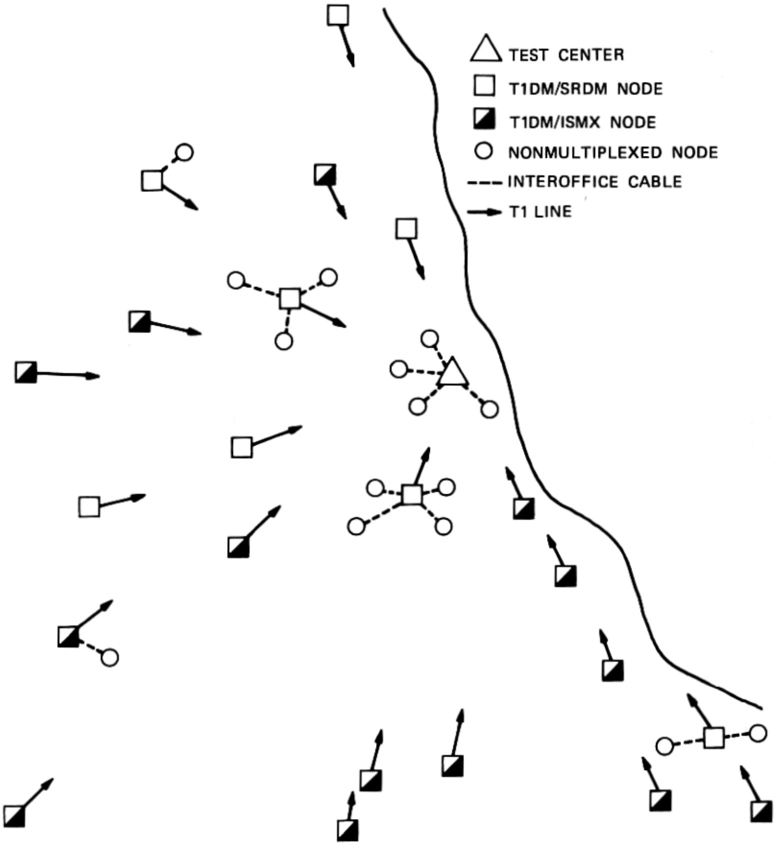

Fig. 6—Typical DSA network.

5.2 Digital voice/data shared transmission

In the example given in the previous section, it was seen that a single T1DM/ISMX multiplexer has capacity for about 80 data circuits. For those node clusters that do not require this capacity, the use of T1WB4s may be advantageous. The T1WB4/ISMX combination can multiplex about 40 circuits with the 50-20-20-10 speed mix assumed above, so it will have sufficient capacity for many small node clusters during the early network implementation. A potential economic advantage of using the T1WB4, however, is that the T1 line can be shared with encoded voice channels, and the line may thus be used more efficiently. For example, if a single T1WB4 is used at a node, at least 12 encoded voice channels can also be carried on the T1 line connecting the node to the test center location. If only 20 subrate data circuits were served by the T1WB4, as many as 20 voice channels would be available.

The T1WB4/ISMX multiplexing combination is generally less costly than T1DM multiplexing for nodes in which the T1WB4 has sufficient capacity for the demand over a reasonable time period. The T1WB4 is thus used at nodes where the data circuit demand for 64-kb/s channels is not expected to exceed 12 for a period of two to three years. The T1WB4/SRDM multiplexing combination is not used.

5.3 Multiplex chaining and hubbing

A third possibility for reducing the cost of the DSA network is to reduce the T1 line mileage required. Two methods of doing this are indicated in Fig. 7.

T1WB4 multiplexers, being three-port devices, can be "chained" on a single T1 line. The T1WB4 at node B in Fig. 7 receives the T1

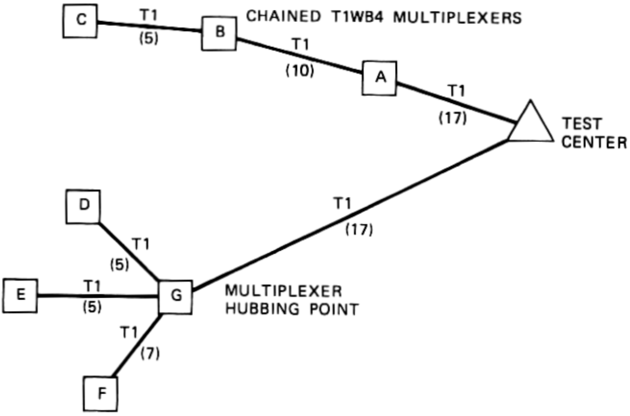

Fig. 7—Digital line sharing.

bit stream from node C, adds some number (less than 12) of 64-kb/s multiplexed channels to the bit stream, and transmits the T1 bit stream on to node A. The T1WB4 at node A operates in the same manner. Although the number of 64-kb/s channels added at any one node is less than or equal to 12, up to 24 channels on the T1 line can be used by the entire set of nodes on the chain. However, if more than 12 channels are used, two T1WB4 multiplexers are required to terminate the chain at the test center.

It is clear that the T1 line mileage required to chain nodes A, B, and C to the test center is less than that required to connect each node individually to the test center. (It is, of course, possible to construct node configurations in which this would not be true, but the statement holds in most practical situations.) Further, a chain of n nodes reduces the number of T1WB4s required at the test center from n to either 1 or 2, depending on whether more than twelve 64-kb/s channels are required in the chain.

It should be noted that the above advantages of chaining are expressed in terms of multiplexer and transmission savings. Maintenance, administration, and system reliability considerations impose a limit of three on the maximum number of links in a chained configuration. Also, voice sharing and chaining cannot be used simultaneously on the same T1 line.

Another method for saving T1 line mileage is also shown in Fig. 7. This involves the creation of a hubbing point at a node very similar to the hub at the test center location. In the figure, nodes D, E, and F are served over T1 lines to node G, where T1 signals are demultiplexed to subrate levels, regrouped, and remultiplexed onto one or more T1 lines between node G and the test center. Although T1 line mileage can be saved in most practical situations, node multiplexing costs are usually increased because of the back-to-back multiplexers required at the hubbing node.

In general, therefore, hubbing nodes are only established at locations anticipated as future test center locations.

5.4 Example

An example DSA network for a large metropolitan area is given in Fig. 8. The plan indicates how multiplexing flexibility can be used in many ways to plan an efficient DSA network.

VI. SUMMARY

The objective of the DDS network planning described in this paper has been to define methods to evaluate trade-offs between node multiplexing requirements and transmission efficiency. In the long-

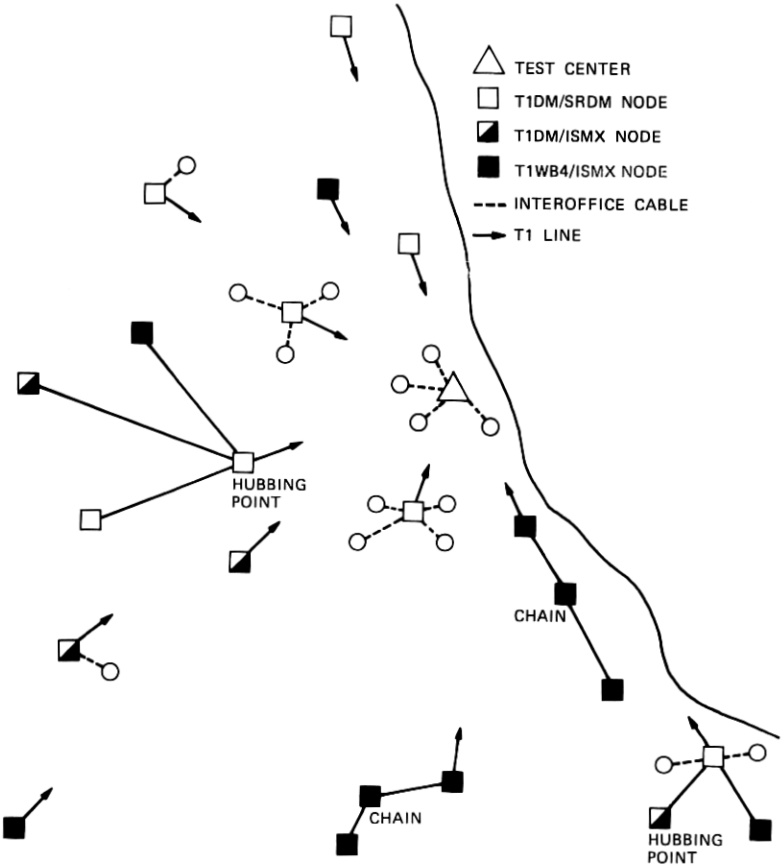

Fig. 8—Typical DSA network for metropolitan area.

haul network, this has led to definition of a three-level bit-stream routing hierarchy, express routing of high-fill bit streams, and techniques for applying these bit streams to least-cost transmission facilities. In metropolitan area networks, emphasis has been placed on first reducing node multiplexing costs by tailoring the multiplexer arrangements to anticipated demand, and then by the use of voice sharing, chaining, or hubbing techniques to increase the efficient usage of T1 lines.

VII. ACKNOWLEDGMENTS

The network design methods described in this paper resulted from close cooperation by a number of AT&T, AT&T Long Lines, and Bell Laboratories representatives. Major contributors have been

R. J. Blackburn, F. J. Ferrantelli, and P. E. Muench of AT&T, K. T. Davis, R. V. Maynard, D. E. Wenski, and G. Squitieri of Long Lines, and W. G. Heffron, B. J. Lifchus, J. J. Mansell, R. E. Reid, E. J. Rodriguez, and H. A. Sunkenberg of Bell Laboratories.

REFERENCE

1. E. W. Dijkstra, "A Note on Two Problems in Connexion with Graphs," Numerische Mathematik, *1*, No. 1 (1959), pp. 269–271.