No. 1 ESS Switching Network Plan

By A. FEINER and W. S. HAYWARD

(Manuscript received January 13, 1964)

An eight-stage space division switching network with ferreed crosspoints was adopted for No. 1 ESS. It has a low crosspoint count and is adaptable to a wide range of office sizes and traffic parameters. This article discusses the network topology, control philosophy, and traffic aspects.

I. INTRODUCTION

This article presents the topology, traffic properties and control of the switching network for No. 1 ESS. Companion articles describe the physical implementation of the network[1] and its control.[2] In the near future, an article dealing with the program control will be published.

To assist the reader with unfamiliar terms, a brief list of definitions is given in the Appendix.

In connection with the plans to develop an electronic central office, the problem of switching network design has received much attention. Due to the many, often conflicting, requirements and possible choices of technology and geometry, the synthesis process requires a fair measure of that ill-defined catalyst commonly referred to as "intuition." The invention and the successful development of the ferreed[3,4] provided the technological solution that resolved the early difficulties of all-electronic networks. The ferreeds provide a metallic transmission path while retaining high switching speed.

Among the early recognized requirements was the desire to use the switching network not only for the obvious function of interconnecting lines and trunks, but also for all link functions — connections between signal transmitters and trunks, ringing circuits and lines, etc. The reasons behind this requirement were to simplify trunk circuits, to eliminate the problem of engineering and administering several different networks, to simplify control, to provide the connecting function at high efficiency and to provide full freedom to associate trunks with all types of signaling circuits.

The field of application envisaged for No. 1 ESS encompasses offices from just a few thousands to many tens of thousands of lines, with line occupancies and ratios of intraoffice to interoffice traffic highly variable from one office to the next. From the standpoint of manufacture and engineering, it is important that this wide range of requirements be met with a single, standard, but adequately flexible, network plan rather than with numerous custom-tailored solutions.

This approach, however, leads to a compromise and results in some loss of efficiency at the extremes of the parameter range.

The total cost of the network can be viewed as consisting of three main parts:

(a) the cost of network crosspoints

(b) the cost of equipment directly associated with network controls (this usually grows with the number of links), and

(c) the proportionate share of central processor cost applicable to the handling of the network.

Freedom from standard matrix sizes embedded in the electromechanical technology (such as crossbar switches) and promise of more subtle network control methods opened the way for considering networks with a larger number of switching stages containing, perhaps, switches of smaller dimension. It was felt that this would permit attaining large network sizes, while retaining a low crosspoint count per line.

The network plan developed for the Morris, Illinois, trial exchange,[5] largely upon ideas of Mr. C. E. Brooks, underwent much scrutiny and served as a point of departure. It was felt that a further crosspoint saving could be made by changing the topology; more importantly, an improvement in growth characteristics was sought.

Independently from studies of central office networks, an exploration of suitable remote concentrator arrangements was being carried out at that time. It yielded the two-stage concentrating configuration shown in Fig. 1.

The first stage of the concentrator is formed by four 16×8 switches in which each of the sixteen inputs has access to only four output links. The placement of crosspoints is identical to the position of diodes in a binary to one-out-of-sixteen diode translator. As a result, every input has access to each of the four pairs of output links. Since these are distributed in pairs over the four-output 8×4 switches, the resulting configuration provides access from every input to every output.

The numerical elegance of this pattern, the economy of crosspoints, and certain other properties (such as the fact that the first-stage switch maps into a full 8×8 switch) made a full-scale investigation appear worthwhile.

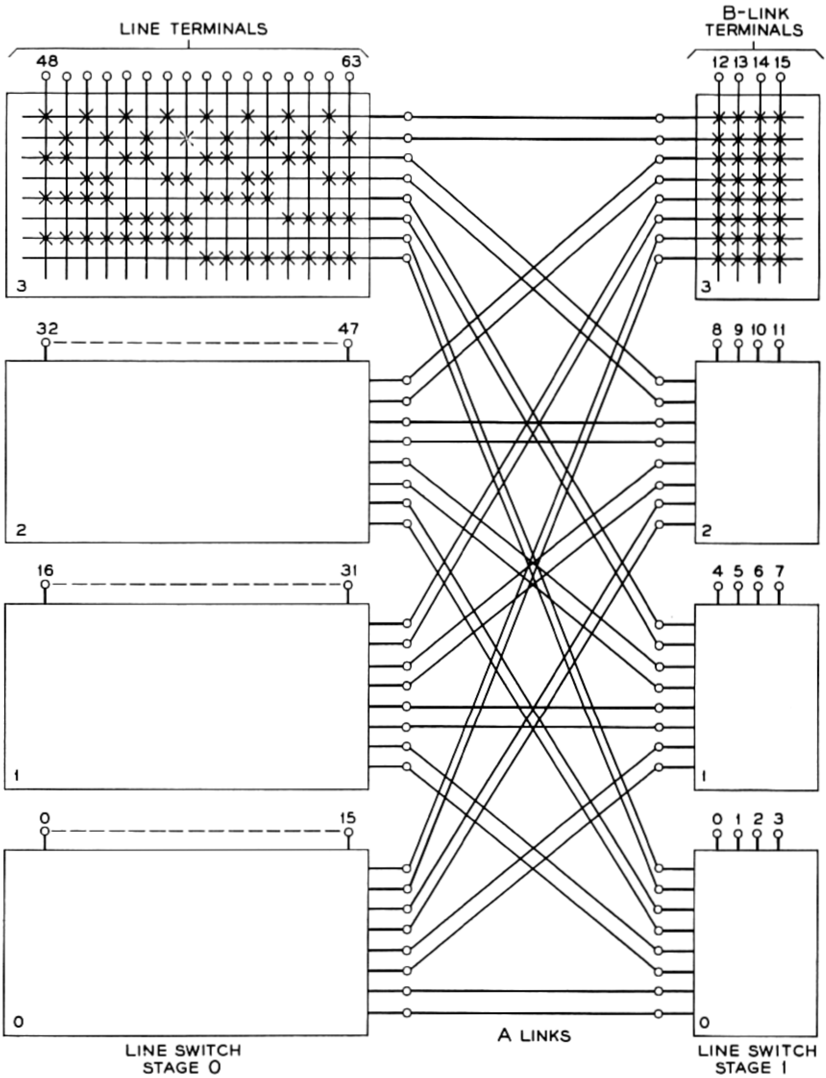

Fig. 1 — A 64-to-16 two-stage concentrator arrangement.

The performance capabilities of this concentrator were studied with the help of a computer simulation; the results will be discussed later in this article.

The attempt to solve the problem for a concentration ratio of four originated in the knowledge that the average occupancy of a line in the

Bell System is about 0.1 erlang and that most economical switching networks investigated at that time were capable of internal link occupancies of 0.4 erlang.

II. THE OCTAL NETWORK

2.1 *Over-All Plan*

The binary nature of the ESS control language led to adoption of switch and grid sizes characterized by numbers of terminals that are powers of two to realize translation and control economies. The choice of switch size was made on the basis of studies of physical design, control cost and number of switches needed to meet objective size and traffic capacity. Of the binary sizes, 4 × 4, 8 × 8 and 16 × 16 were obvious contenders; 8 × 8 was chosen. Considerations of access and blocking in the largest network size led to the adoption of a network with eight stages of switching.

Topologically, the network consists of four-stage groupings of which there exist two types — the line link networks and the trunk link networks. Connecting these subnetworks among one another are junctor groups provided in a pattern consistent with the specific size and traffic character of a given office.

As the name implies, the subscriber lines connect to the line link network. Two basic sizes exist for the line link network. One, with a concentration ratio of 4:1 in the first two stages of switching, provides terminations for 4096 subscribers; the other, developed with higher traffic loads in mind, has 2:1 concentration and provides terminations for 2048 inputs. A constant number of 1024 junctor terminals characterizes all (fully equipped) link networks.

Trunks and service circuits connect to the trunk link networks; these have the basic size of 1024 trunk terminals.

Fig. 2 shows an example of the network in an office with approximately 8000 subscribers and 1000 trunk and service circuits. A novel feature of the octal network is the method of handling intraoffice calls; these are routed on direct intraoffice junctors that link the line link networks among one another and to themselves and bypass the trunk link networks. The intraoffice junctors contain the circuitry to apply battery and to supervise both subscribers.

2.2 *The Line Link Networks*

The first two stages of the line link network contain a concentrator arrangement, shown in Fig. 3(a) for the 4:1 ratio and Fig. 3(b) for the 2:1

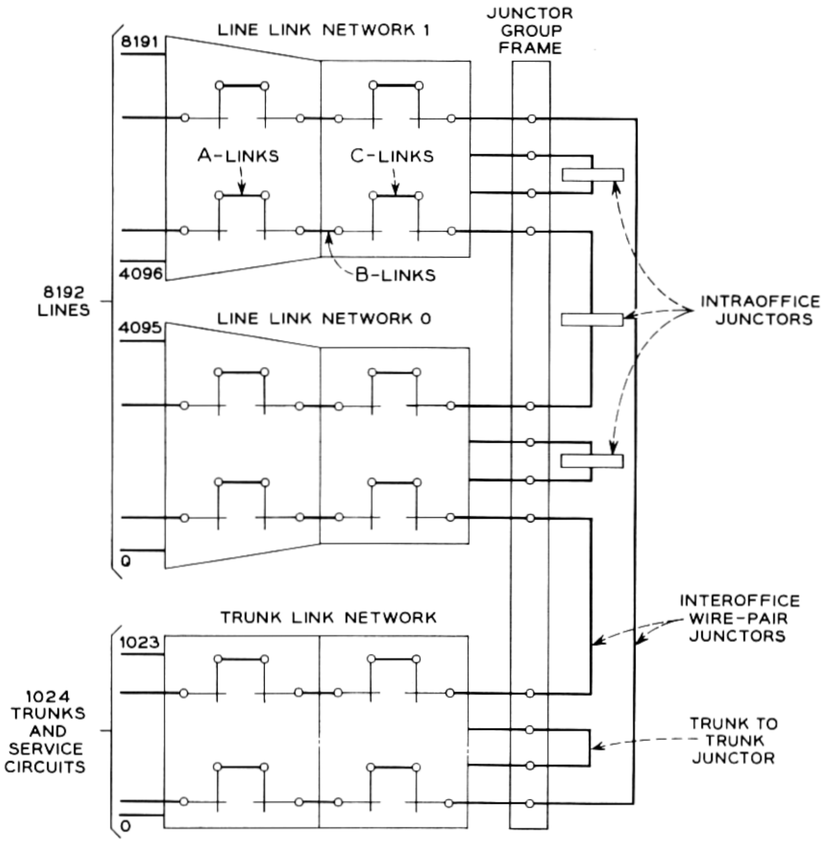

Fig. 2 — An example of the network for an office with 8000 lines and 1000 trunk and service circuits.

ratio. The first is a slight modification of the previously described concentrator; the change in the first-stage switch pattern was found to sacrifice little traffic carrying capability and simplified the internal structure of the ferreed switch. The third and fourth switching stages are formed from 8 × 8 switches in both types of line link networks, organized into 16 grids (see Fig. 4). This configuration provides every line with full access to the 1024 junctors. It is convenient when dealing with grid networks of this type to express the access within the network as the product of the individual switch access numbers

$$A_{LLN} = a_1 \cdot a_2 \cdot a_3 \cdot a_4$$

and with $a_1 = a_2 = 4$ and $a_3 = a_4 = 8$

$$A_{LLN} = 1024.$$

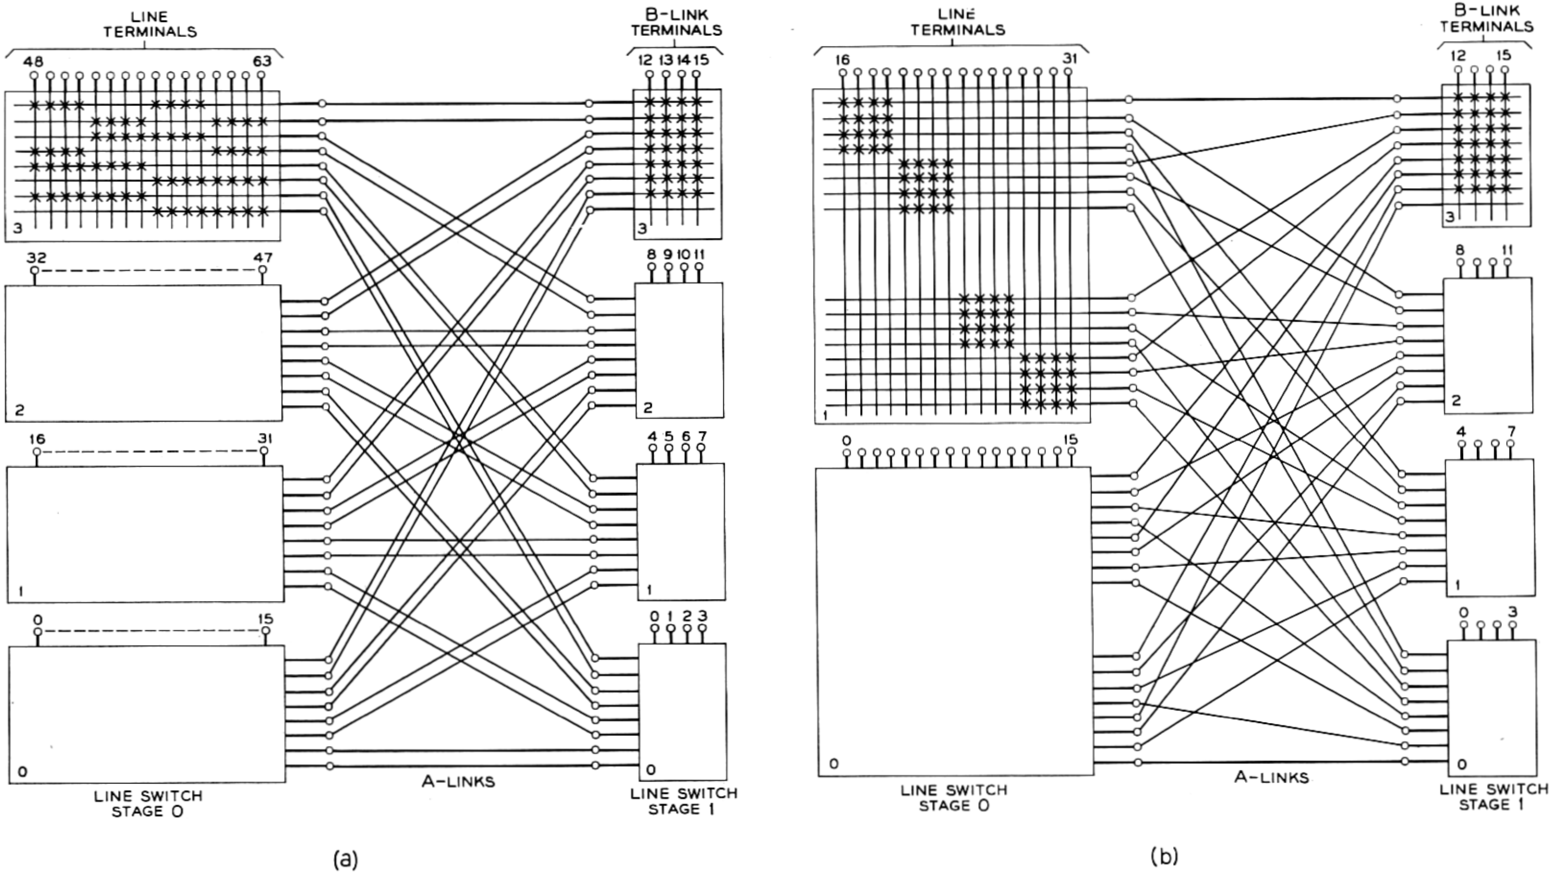

Fig. 3 — (a) A modification of the 4:1 concentrator of Fig. 1 for No. 1 ESS network. (b) The 2:1 concentrator connection pattern.

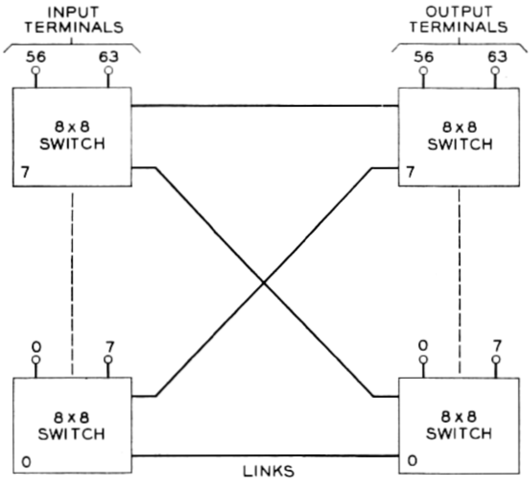

Fig. 4— The 8 × 8 grid.

This provides a quick verification of the full-access nature of the line link network; needless to say, proper link distribution between the concentrating stages and the 8 × 8 grids must be observed. In this case there must be one B link between every concentrator and every grid.

The equipment design of the network is discussed fully in other articles.[1,2] It will suffice here to say that the line link network contains two types of frames — the line switch frames, into which are packaged 16 concentrators, and the junctor switch frames containing eight octal grids each. In addition to the crosspoint arrays, the line switch frames house the equipment that constitutes the line circuits, namely the line scanner and the ferreed devices for removing the current-sensing ferrod element from the line after a service request has been detected and registered. The junctor switch frames also contain additional ferreeds, one per junctor, that provide test access into established network connections. Fig. 5(a) shows the composition of a line link network in terms of these equipment units.

All of these equipment frames contain their own duplicated control circuits. These circuits receive and translate orders from the central processor and perform all the other functions that lead to their execution. Typically, an order calls for the closure of a specified path through the two switching stages contained in the frame.

The two duplicated control units work independently of each other; when both are functioning properly, each restricts its activity to its assigned half of the network contained in the frame.

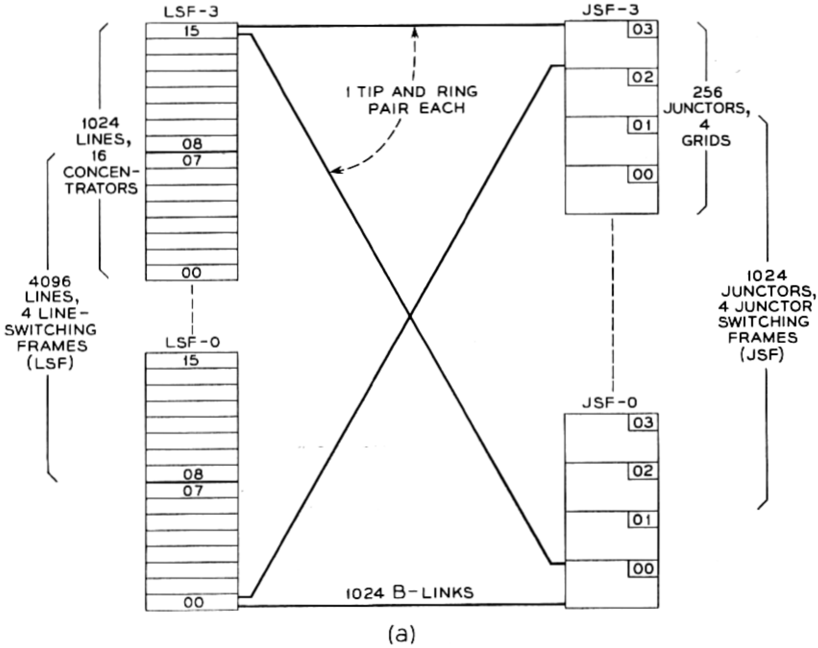

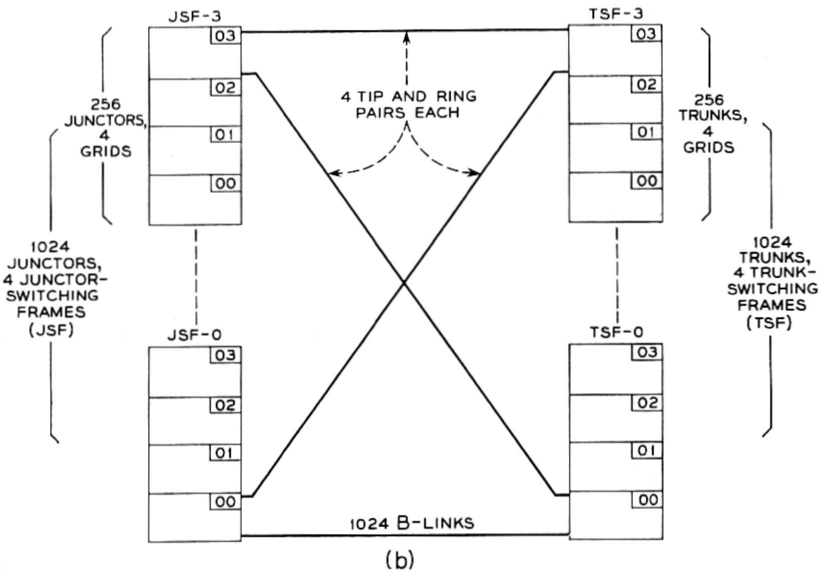

Fig. 5 — (a) Line link network (with 4:1 concentration). (b) Trunk link network.

2.3 *The Trunk Link Network*

When the trunk link network was first studied, it was proposed to construct it as a three-stage network containing a 16 × 16 switch at the trunk side followed by two stages of 8 × 8 switches. This naturally resulted in the convenient full-access subnetwork of 1024 inputs and 1024 junctor outputs (16 × 8 × 8 = 1024). Further investigation has shown, however, that at the slight control complication of introducing another stage, a substantial gain in traffic carrying capacity could be realized with the same number of crosspoints per terminal by going to four stages of 8 × 8 switches.

Since the trunk link network access is thus increased to $A_{TLN} = 8^4 = 4096$, if the same size of the trunk link network of 1024 × 1024 is retained, a multiple access from trunks to junctors is obtained with every trunk capable of reaching every junctor by four different paths. This reduces considerably the trunk-to-junctor blocking despite the addition of a stage of switching.

Two types of equipment frames are contained in the trunk link network. One of them is the junctor switch frame that also serves as a building block for the line link network. The other, the trunk switch frame, contains the same number of octal grids but has none of the test access provisions of the junctor switch frame. Fig. 5(b) shows the composition of a trunk link network.

2.4 *Interconnection within Line Networks*

Fig. 6 gives a three-dimensional view of a full line link network, showing in skeleton form the way in which frames are connected together. Concentrators are shown as horizontal planes; octal grids are shown as vertical planes. The connections between line switch frames and junctor switch frames, B links, are shown only for the edge of the link network. This diagram gives a picture of how all lines have access to the 16 B links of their concentrators and of how each B link connects to a different octal grid in the junctor switch frame. The grids in turn give each B link access to 64 junctors. Fig. 7 gives a similar representation of a full trunk link network. Pictorial representation is more difficult here. Each trunk has access to 64 B links which are treated as four groups of 16 links. Each group of 16 covers the 16 octal grids of the junctor switch frames.

2.5 *Partial Network*

The size of a link network is large. Equipment frame sizes, on the other hand, have been so chosen that the incremental cost of buying equipment

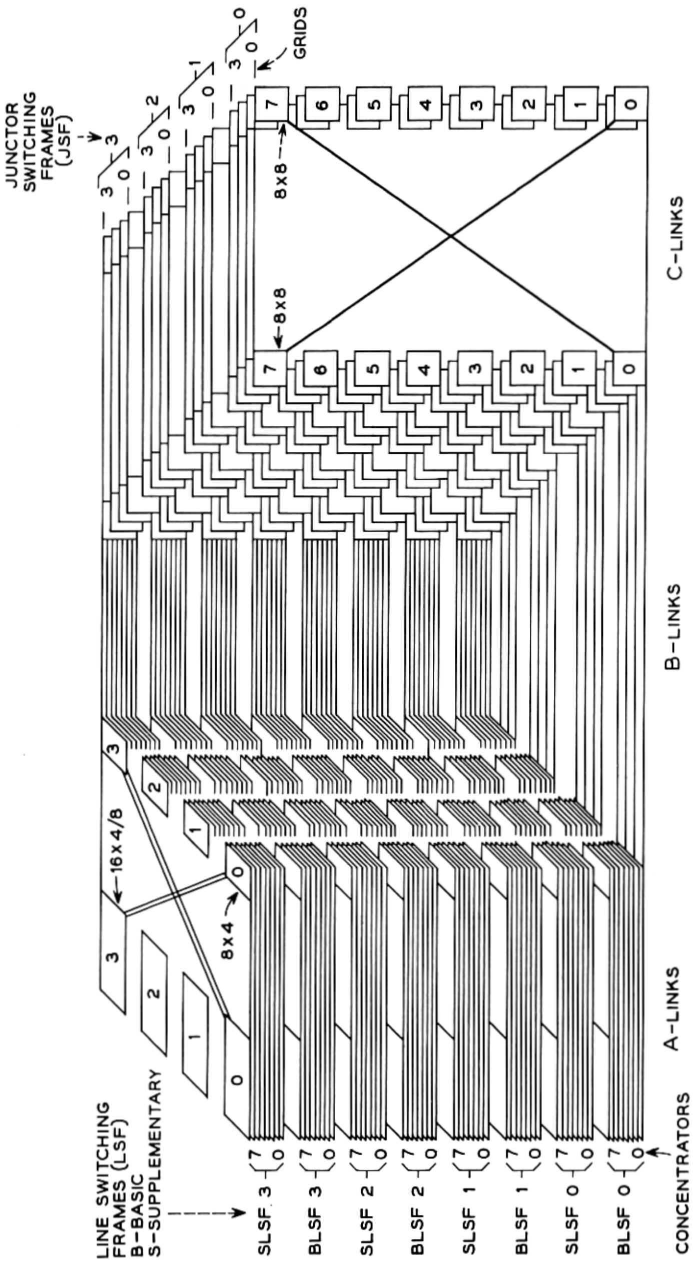

Fig. 6 — Schematic diagram of line link network.

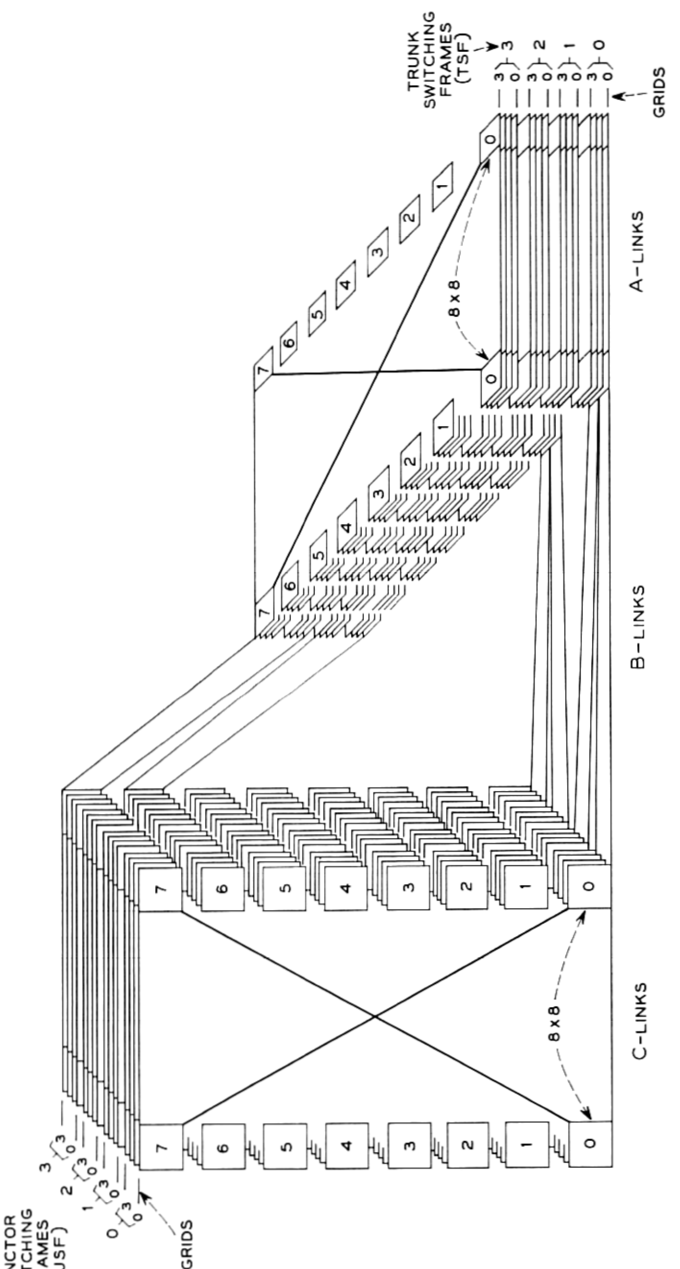

Fig. 7 — Schematic diagram of trunk link network.

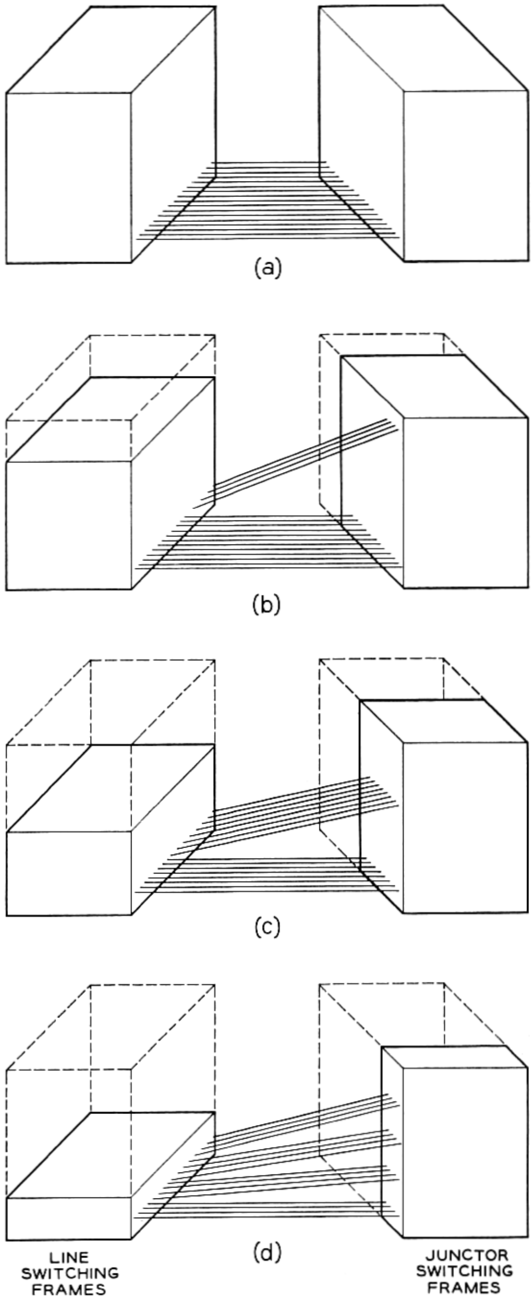

Fig. 8 — Schematic diagram illustrating the partial equipping of line link networks.

can be kept within reason. In the frame design, a compromise has been made between the savings resulting from manufacturing equipment in large units and the excess amount of equipment that may be purchased in each installation because of the necessity of buying integral numbers of frames. As mentioned previously, line switch frames provide terminals for 1024 lines (512 for the 2:1 line concentrator), junctor switch frames for 256 junctors and trunk switch frames for 256 trunks. In order to equip a network with fewer than its full complement of frames, a special wiring plan must be used. Line or trunk switch frames are easily omitted, since they take with them their B link traffic. To omit a junctor switch frame, however, it is necessary to reassign some of the B links which carry traffic from the line or trunk switch frames. Fig. 8 shows in a simplified form based on Fig. 6 how a group of 16 B links from a line concentrator is reassigned for quarter, half, three-quarter and full line link networks. Similar patterns are used for the trunk link networks, although they are somewhat more complicated by the large number of parallel paths. Partial equipping in either link network increases the number of parallel paths while reducing the number of junctors available for connections to other link networks. The two offset each other to a certain extent, so that partial link networks can be traffic loaded almost as efficiently as full link networks.

2.6 *Network Sizes*

Up to now a line link network of 2048 or 4096 lines and a trunk link network of 1024 trunks have been described. If the line and trunk average usages happen to be just right, these link networks carry traffic with high efficiency. If, however, the load per line or trunk is too high, not all terminals can be assigned. On the other hand, if the terminal load is too low, some of the switching equipment is wasted. To reduce the waste in the latter case, wiring patterns have been evolved which provide for up to twice as many line or trunk switch frames as make up the basic network. Thus up to 8192 lines may appear on a line link network and up to 2048 trunks on a trunk link network. This higher concentration ratio from terminal to junctor is achieved by multipling B links at the third stage of the network. Fig. 9 shows the effect on crosspoint requirements in a line link network with an upper traffic bound of 0.4 occupancy on its links. (Cutoff and test access crosspoints are included.) Unfortunately, except in the maximum size, one cannot merely add new links to the basic links without rearrangement. To do so would result in the added equipment being served only by shared links, while the basic equipment would be served by a mixture of shared and private links. The resulting

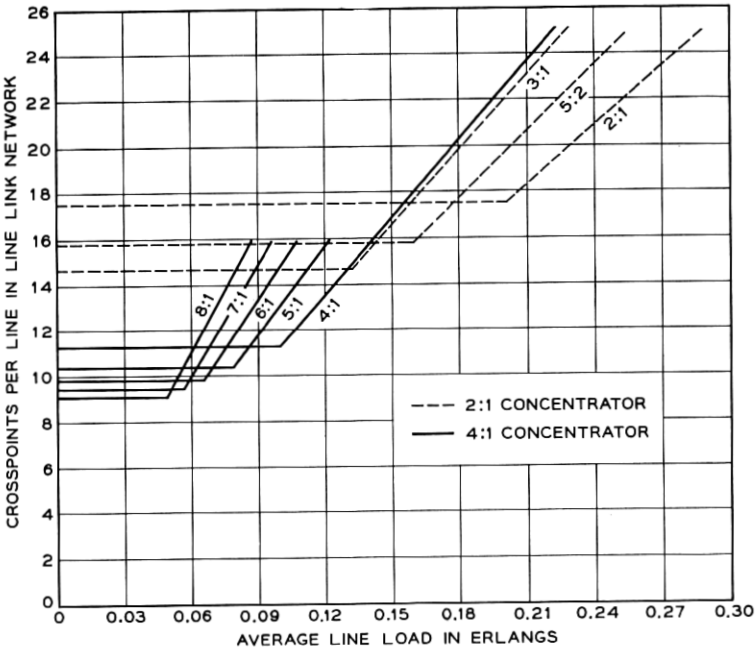

Fig. 9 — Crosspoints per line as a function of line usage when an upper bound of 0.4 occupancy is selected.

uneven service is unacceptable; instead, the patterns are arranged to provide as nearly equal sharing as possible. In choosing the patterns, attention was given to reducing the number of link reassignments that would be necessary if the average load per line or trunk should change significantly. Patterns for partial equipping are unaffected by the choice of concentration ratio.

Central offices with heavy PBX development may have line usage as much as twice the usual average. A blank terminal on a line switch represents not only wasted crosspoints, but a wasted scan point, cutoff contact and main distribution frame appearance. As the demand appeared high enough, the second design of line switch frame was made in which 512 lines reach 256 B-links. By similar multipling of B-links this frame can fill the gap between average usage and double usage on lines.

2.7 *Junctor Patterns for Growth*

The link network junctor terminals are cabled to plugs and jacks to make possible orderly transitions from one size office to another, both at the time new equipment is installed and afterwards to take care of changing traffic patterns. Fig. 9 shows how the 64 subgroups of a network are cabled to the junctor grouping frame. A subgroup contains 16 junctors, each junctor from a different octal grid. Half of the subgroups are jack ended and half plug ended — one jack or plug per subgroup. Interconnections are made only within the shelves shown on Fig. 10. Plugs and jacks always interconnect different numbered switches to insure two different sets of C links on connections originating and terminating within a link network.

The junctor grouping frame is also used to insert junctor circuits in

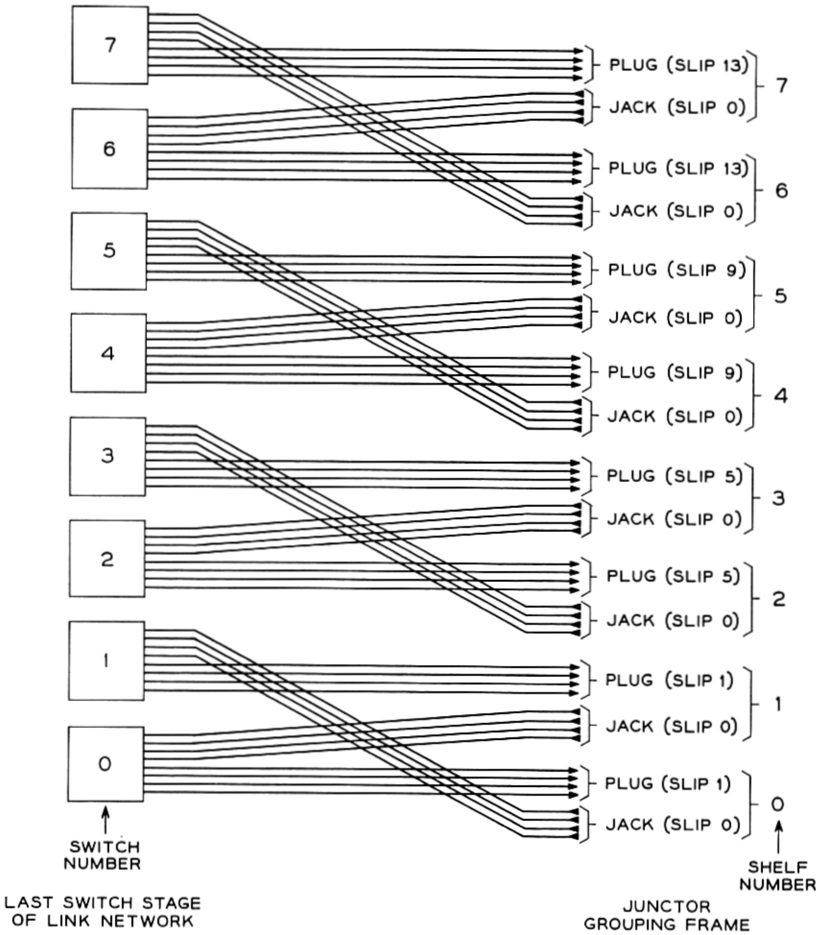

Fig. 10 — Junctor group assignments.

line-to-line junctor subgroups. One side of each circuit terminates on a plug, the other on a jack. Connecting a line-to-line subgroup consists of inserting a line link network plug into a junctor circuit jack, then inserting the corresponding junctor circuit plug into a line link network jack.

It is convenient to use the linear-graph representation of the network in discussing the next aspect of junctor connections. Fig. 11 shows the links available for paths between two lines with a subgroup connected plug-to-jack on shelf one (solid) and shelf three (dashed). The junctors are seen to be "slipped" by one and five terminals, respectively. A slip of at least one is necessary for making a connection between lines assigned to the same concentrator — a slip of zero would provide no path for intraconcentrator calls because the same B link would be required twice on any path. A slip of at least four is necessary for completing calls between lines on the same line switch. The additional choice of paths provided by giving a second set of junctors a slip which differs from the slip of the first by at least 4 gives many more opportunities of finding paths than a simple parallel choice. The second set aligns both the A links and B links of the two ends of the connection in different combinations. As noted in Fig. 10, the plugs are wired for slips of 1, 5, 9 and 13. The definition of an actual slip depends on the point of view. The solid wiring of Fig. 11 gives a slip of 1 when viewed from left to right, but a slip of 15 when viewed from right to left. Thus 8 slips (1, 3, 5, 7, 9, 11, 13, 15) are available in the assignment of junctors — governed by the choice of shelf on which junctors are connected and whether the connection is plug-to-jack or jack-to-plug.

When networks are partially equipped with either one, two or three junctor switch frames, the junctor wiring is changed so that full junctor subgroups still appear on a plug or on a jack. As with the B links, the particular choice of wiring is made with the objective of reducing the number of wires that must be moved as the network grows.

III. NETWORK CONTROL

3.1 *Path Searching*

A basic decision in the design of the switching network was that of isolating completely the path searching function from the switch itself. External control circuits cannot determine directly the states of the crosspoints within a switch. In accordance with the general No. 1 ESS approach, the central processor makes all path searches and keeps a continuous record of all pertinent switching information in its temporary

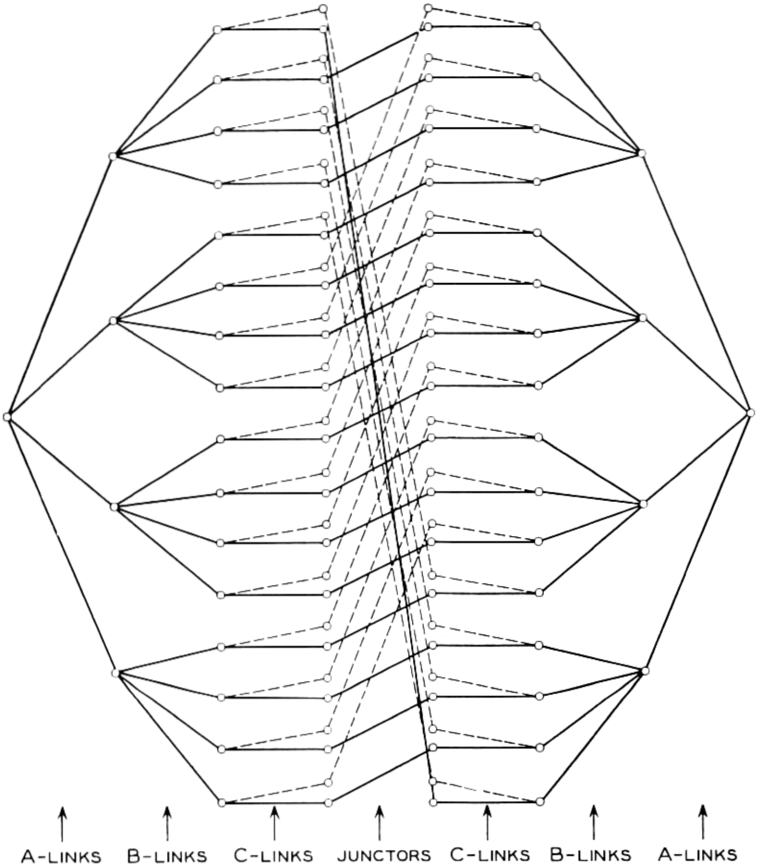

Fig. 11 — Linear graph of line-to-line paths with two junctor subgroups.

memory, the call store. Programs which use this information, either in setting up paths or in releasing them, must keep the network records up to date. Network records are among the most vital of those kept in call store. Their loss would be equivalent to the loss of power in an electromechanical switching network. The memory reliability has been made high. Beyond this, the network control programs have been designed to keep the chance of error low.

The format of the switching network record in call store was chosen with a view toward low processing time in establishing or releasing network connections. The records are somewhat redundant because of this objective but the redundancy also provides additional insurance against memory failure. There are two basic records: "link memory" is provided on a basis of one bit for each link and is used in the path searches. A "0" indicates a busy link; a "1" indicates an idle link. "Path memory" is provided on a basis of one word for each junctor terminal in a line link network and one word for each trunk terminal in a trunk link network. It is used to store data necessary for releasing connections.

Let us consider first the link memory and, for example, the problem of finding a path between two lines. This can be divided into searches for a path from a line to a junctor on the terminating and on the originating link networks and then a search for a commonly accessible junctor. Fig. 12(a) gives a graph of a typical situation showing link status. Fig. 12(b) gives the corresponding link words. The 16 lines terminated on a line switch have access to 8 A links. The bits for these links (and for the 8 links in an adjacent switch) are contained in a single link bit word. Through the A links, each line has access to 16 B links. Again the 16 bits corresponding to these links are contained in a single word. Now a path through the network must be set up through an idle A link and an idle B link. By suitable masking and shifting, the bits corresponding to the 4 eligible A links can be extracted, and from them a 16-bit word can be generated with 4 bits for each A-link bit occupying positions corresponding to each of the eligible B links. The resultant 16-bit word can now be combined using the logical AND central control function, with the B-link bit word to produce a matching word in which 1's represent those idle links with free paths to the particular line. Continuing through the network, each of the B links can reach 8 C links. The link bits corresponding to these C links are so arranged that a bit in a C-link word represents one of the C links accessible to a B link (represented by a bit occupying a corresponding position in a B-link word). A C-link word can therefore be combined with the matching word to test 16 paths one stage further into the network. A similar action on a junctor link bit word results in a matching word indicating all available paths from line to a selected subgroup of 16 junctors. Junctor connections between link networks are always made in integral numbers of these junctor subgroups. Different combinations of link and junctor words can be used to match paths from a line to any of the 64 subgroups of a full link network.

After a matching word is determined between originating line and junctor subgroup, a similar word can be derived for the terminating line and the same subgroup. Assuming that idle paths exist in both words, it is now necessary only to take into account the junctor slip. Because of this slip the bit positions in the matching word for one network will not correspond to those in the matching word of another network. One of

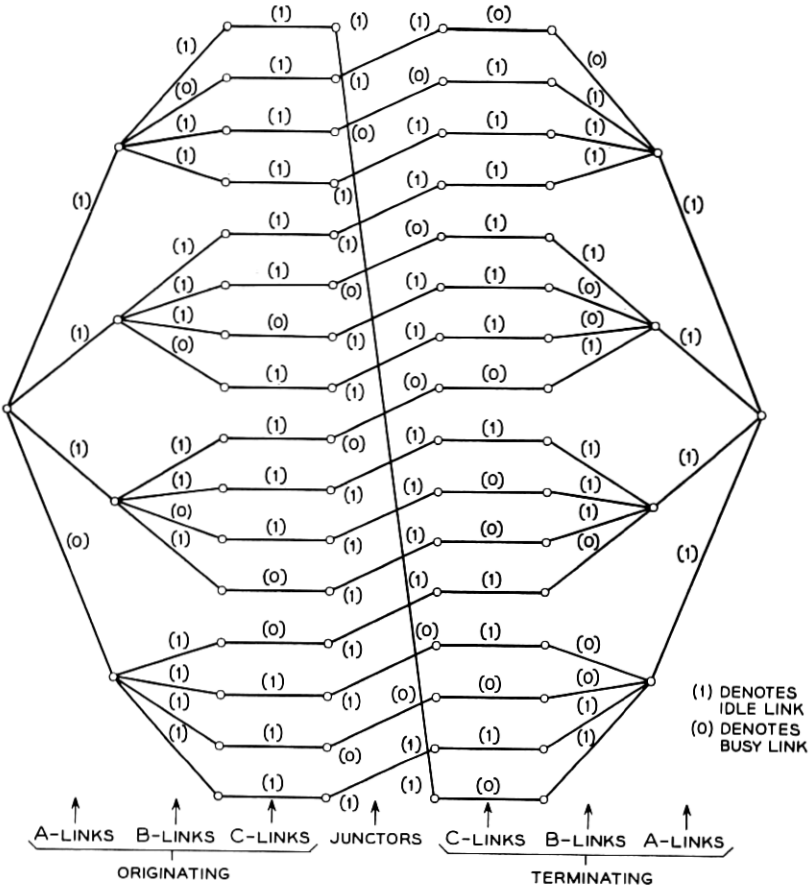

(a)

	ORIGINATING	TERMINATING
A-LINK	1110 1101 1110 0110	1011 1111 1110 1111
EXPANDED A-LINK	1111 1111 1111 0000	1111 1111 1111 1111
B-LINK	1011 1110 1101 1111	0111 1001 1110 0011
A-B	1011 1110 1101 0000	0111 1001 1110 0011
C-LINK	1111 1101 1110 0111	0111 1110 1001 1010
A-B-C	1011 1100 1100 0000	0111 1000 1000 0010
JUNCTOR	1101 1011 0111 1101	1011 0110 1111 1011
A-B-C-J	1001 1000 0100 0000	0011 0000 1000 0010
ROTATED ORIGINATING WORD		0011 0000 1000 0001
MATCHING WORD		0011 0000 1000 0000

(b)

Fig. 12 — (a) An example of the states of the links of a line-to-line connection. (b) Searching a path.

the words is "rotated" to line up the path bits and then the two words can be combined to see if any free path exists. If so, the central control order for finding the leftmost 1 can be used for fast identification of the path. If not, a second junctor subgroup must be tested. Because the above method tests 16 paths at a time and because there is a reasonably high chance of success on the first junctor subgroup chosen, searching time is relatively low. Other concentration ratios and partially equipped link networks require variations on the steps given in this example.

A further complication, not shown above, is the necessity of reusing links which have been in use on a previous section of a call or which are being reserved for an anticipated connection. Thus, in the example, at least the A and B links which were used in the dialing connection to the originating line should be available for the line-to-line connection used. Similarly, the A and B links reserved for the line-to-line connection should be available for ringing connections. Failure to make these links available would drastically reduce network capacity.

Link bits contain sufficient information to hunt an idle path between two network terminals. They do not, however, contain enough information to identify a path for releasing a connection, because they do not indicate which link is connected to which. This information is stored separately in blocks of call store named "path memory words." Line path memory words are assigned in blocks of 1024, with each word corresponding to a junctor terminal of a line link network. Trunk path memory words are assigned in blocks of 256, with each word corresponding to a trunk terminal of a trunk switch frame. A line path memory word contains the identity of the line connected to the associated junctor, while a trunk path memory word contains the identity of the junctor connected to the associated trunk. Additional bits in these words identify the path used when more than one path is possible between the junctor and line or trunk.

The correspondence between path memory word and network terminal was chosen to simplify most of the programs which change link bits when a path is released. Thus when a trunk indicates that a network path is no longer needed, translation from trunk identity to network termination serves also to locate the trunk path memory word. This in turn identifies the line-to-trunk junctor number which serves to locate the line path memory word. Line-to-line connections are traced in a similar fashion, but additional information is needed on trunk-to-trunk connections because only half of the path can be traced from the trunk path memory word, and there is no path memory associated with trunk junctor terminations. On trunk-to-trunk connections, therefore, additional path memory is provided in a register associated with the call.

3.2 *Network Actions*

Searching for a path is only part of the network control job. Orders must be issued to the switch frame controllers to close the specified crosspoints, to signal distributors to open or close relays in trunks, and to scanners to verify that all orders have been properly executed. Because the time restrictions on the network are given in tens or hundreds of milliseconds, relays have been used in both switch frame and signal distributor controllers. To match the high-speed central control with the slower-speed relay circuitry, a cyclic method of controller operation has been adopted. Orders are issued in batches at about 25-millisecond intervals. The program which sends out orders keeps its own record of those controllers it has called and will not send two orders to the same controller in one batch. This record is kept in call store, where it can be interrogated at central control speed without waiting for the slow response time of the controllers themselves.

A new batch of orders is sent when four full five-millisecond intervals have passed after the last order of the previous batch; since the interval in which the batch was sent cannot be counted, the expected time between batches is 25 milliseconds. To meet this method of timing, the controllers must finish their work in under 20 milliseconds. At extreme traffic loads, the time taken to issue the orders, when added to other essential work, may stretch many cycles out to 30 milliseconds.

Network actions are controlled in two parts. The first program prepares a list which is placed in a call store area named a "peripheral order buffer." This list contains instructions for the proper issuing of controller orders and is held until the connection is set up. The second program scans the controllers at the start of a network cycle to check that all are ready to receive new orders. It then works through the list, sending out as many orders as it can. When it reaches a point at which it cannot proceed further, either because of a planned delay or because it finds an order to a controller that has already received an order in that cycle, it proceeds to the other buffers which have waiting work. If it finishes a list, it reports back to the call processing program which had requested the connection.

Consider, for example, the task of setting up a path between an incoming trunk and a line. At the time this connection is to be made, the calling trunk is connected to a ringing tone trunk and the called line is connected to a ringing trunk. The functions to be performed are, in order:

(1) open cut-through relay in ringing trunk*
(2) open cut-through relay in ringing tone trunk*
(3) open cut-through relay in incoming trunk
(4) delay for one network cycle to allow relays to open
(5) close second-stage crosspoint in line concentrator
(6) close two crosspoints in trunk switch frame
(7) close two crosspoints in trunk junctor switch frame
(8) close two crosspoints and test access crosspoint in line junctor switch frame; make false cross and ground test
(9) wait one cycle for test to be completed
(10) open test access point in line junctor switch frame
(11) close last crosspoint in line concentrator
(12) wait one cycle for crosspoints to act
(13) close cut-through relay in incoming trunk
(14) wait two cycles for relay action and for transient decay
(15) scan for line current at the incoming trunk to verify continuity of connection
(16) report back to originating program.

The network control program places the list of orders for this sequence in a peripheral order buffer. At 25-millisecond intervals the list processing program will work as far through this list as it can get. Assuming no delays from busy controllers, the first delay will occur at step (4). A pointer will be left at step (5) so that on the next cycle the list processing can be taken up again through step (9), and so forth, until the path is verified.

IV. TRAFFIC PERFORMANCE

As indicated earlier, the large number of meshed paths in the network makes possible a high efficiency of the network links for objective levels of blocking. In actual use this high efficiency must be weighed carefully, since, as with most traffic carrying facilities, an increase in efficiency at a given service point results in more rapid decline in service at loads beyond this point. The evaluation of a network design includes an examination of a range of loads to make sure that operating points can be found with sufficient range of good service. The final choice of the operating point is an administrative responsibility of the operating telephone company.

* These are the only relay actions required to release the network path. The other parts of the path will be used as needed without interference because of the differential excitation of the ferreed switch.[3,4]

The performance of this network is conveniently broken into two parts, that of the line concentrator by itself and that of the full network in its various uses.

4.1 *Line Concentrator Characteristics*

Fig. 13 shows a load-service curve for the concentrator based on simulation results. This simulation was applied to the concentrator alone; it consisted of repeated offerings of simulated calls to a concentrator modeled in a general-purpose digital computer. The data for Fig. 13 were collected for a concentrator with all lines assumed to have equal calling rates and usage, and are compared with data for a similar theoretical group of 40 lines with full access to 10 links. While this kind of curve is of interest for general evaluation, more concern is felt about the situations which will be found in the field, where equal line usage is the exception rather than the rule. In particular, concern is felt for situations where, through chance, several high-usage lines are assigned to the same switch. With purely random assignments this chance can be shown to be small — it can be made much smaller by adopting assignment practices which tend to spread lines known to be high-usage (such

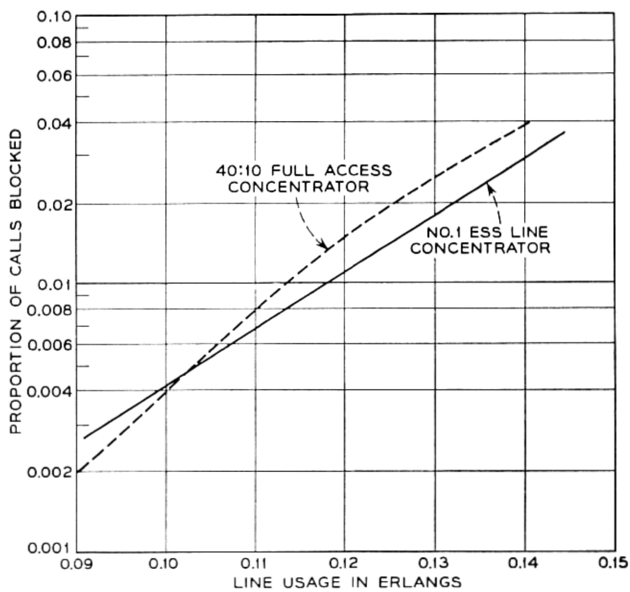

Fig. 13 — Load-service curve of 4:1 concentrator, determined by simulation, compared with theoretical load-service curve for full-access group of 40 lines on 10 links.

as PBX lines) over the concentrator switches. For further simulation, line usages were assigned at random from an exponential distribution. Fig. 14 shows the results of a typical run under these conditions. The average blocking (measured as the number of blocked attempts divided by the total number of attempts) is materially less than the blocking of Fig. 13. On the other hand, the 16 lines on the switch with the highest load experienced a slightly higher blocking.

The drop in average blocking with a wide spread of line usage can be explained as the double effect of the narrowing of the load variance accompanying a variation in source loads and a reserved path effect whereby a high calling rate line may place a new call over the path it has just abandoned before a call from another lower calling rate line can seize it. The average load within the groups of 16 lines is not as closely correlated to the service encountered by the group as might be expected. Closer examination of the groups indicates that the spread of line loads within the group itself is a strong factor. Allowing for the difficulty of predicting the exact performance of each concentrator, it appears safe to use the equal-usage curve of Fig. 13 as a basis for deriving engineering procedures.

In addition to the blocking of the concentrator, the delay given to those originating calls which are blocked is of prime importance. Because new trials are made through the rest of the network to find customer dial pulse receivers, the main network source of this delay is in the con-

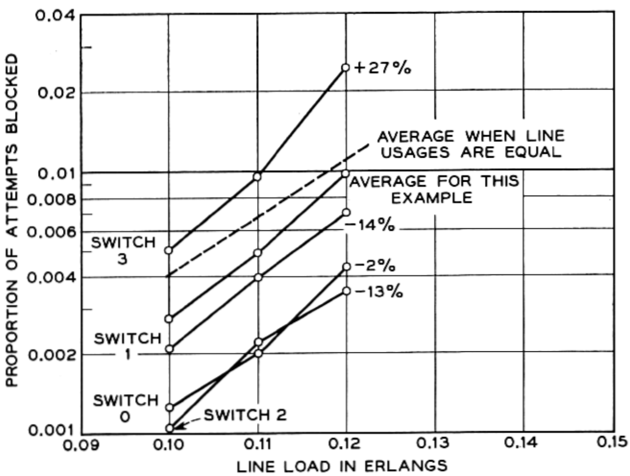

Fig. 14 — Example of effect of switch unbalance on 4:1 line concentrator.

centrator itself. At the completion of dialing, the path through the concentrator which was used for dialing is available for reuse on the ringing and talking connections. In order to evaluate the delay generated by the concentrator, the simulation program was modified to treat all calls as originating calls and to hold delayed calls until they were served. This made possible a direct comparison with the theoretical performance of a full-access group.[6] Fig. 15 shows the results of a simulation run with 0.15 erlang per line, which resulted in slightly over 0.04 of the calls being blocked. The distribution is shown only for the calls which were blocked and is compared to similar theoretical data with a full-access group of 40 lines reaching 10 links. In both cases it is apparent that these groups must be run normally at low losses because dial tone delays, when they occur from this cause, are long. The No. 1 ESS line concentrator with 6 crosspoints per line compares well with the full-access concentrator with 10 crosspoints per line.

The 2:1 concentrator has not been studied in as much detail as the 4:1 concentrator. Not only will it have much lower blocking at comparable link loads because of only second-stage concentration, but the smaller

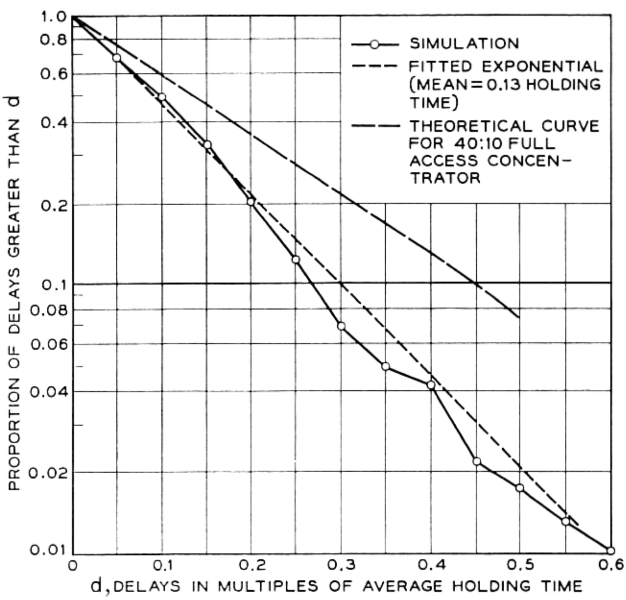

Fig. 15 — Delay distribution of blocked calls, determined by simulation at 0.15 erlang per line, compared with theoretical delay distribution for full-access group of 40 lines on 10 links.

number of lines will introduce a stronger limited-source effect. Under these conditions the network as a whole will be the main source of congestion, and loading of the 2:1 concentrator is not a serious factor.

Because of the delay characteristics of the concentrator, it is expected that the concentrator will be used at low blocking probabilities. If a point such as 0.10 occupancy of the lines is chosen, increases in load of 40 per cent are seen to produce a significant increase in blocking, but this is minor compared to what will happen in the rest of the switching network and interoffice trunks if such high traffic overload occurs.

Fig. 16 shows an estimate of the performance of the switching network

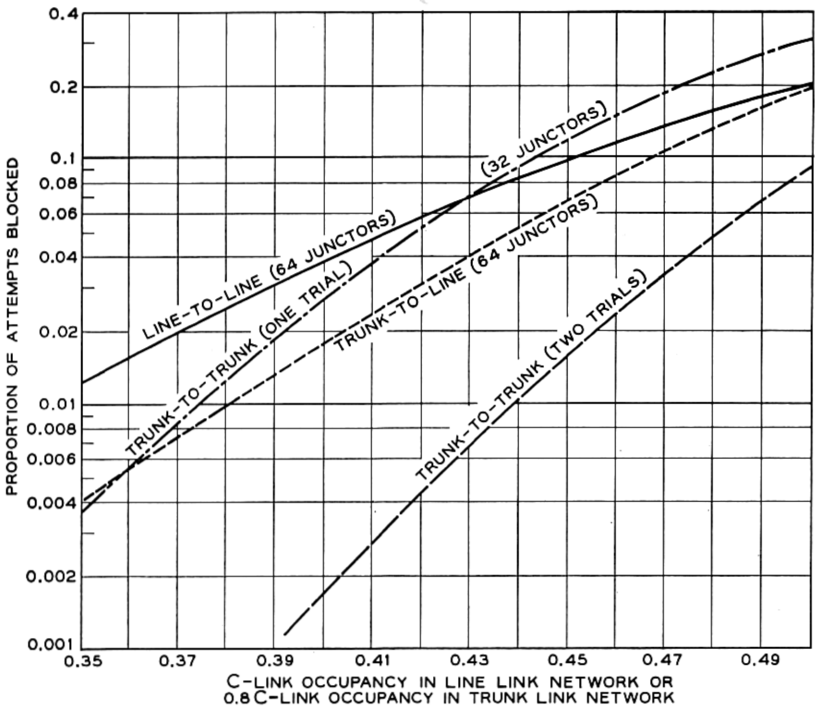

Fig. 16 — Load-service curves of switching network.

(including blocking contributed by the line concentrators). Trunk networks and junctors terminating on trunk networks were assumed to be loaded 20 per cent above the line link networks. All of these data were based on NEASIM* simulations which in turn were verified by a small

* NEASIM[7] is the name given to a computer program designed in Bell Telephone Laboratories. It uses a technique of simulating link states of a linear graph.

number of large simulations by digital computer, in which the network was fully represented. For a large network these full-scale simulations are expensive and generate little data not available by NEASIM techniques.

In actual practice, of course, varying numbers of junctors will be available between link networks; higher efficiencies are available with larger groups. In such cases, the line concentrator may well be the limiting item in loading line link networks of a small office and the intraoffice junctor group the limiting item in a large office. The latter limitation can be overcome by providing intraoffice trunks on the trunk link networks and letting them carry the traffic which overflows small groups of high-efficiency line-to-line junctors.

As with other networks in the past, it is to be expected that initial installations of the network will be over-provided with switches to insure good service until operating experience gives additional data for more precise traffic engineering. Should either the estimates of office traffic or the estimates of service characteristics be significantly in error, the flexibility of junctor assignments and the ability to change the concentration ratio of the link networks offer insurance that efficient operating points can be found.

V. ACKNOWLEDGMENTS

Mr. R. F. Grantges proposed the NEASIM program, which has been of great importance in evaluating network patterns. He has also made frequent contributions to the network plan. Many others have also contributed, in particular, T. N. Lowry, J. G. Kappel, W. C. Jones, C. Klingman and R. A. Rosenthal.

APPENDIX

Definition of Terms

Switching network — that part of a switching system that establishes transmission paths between pairs of terminals.

Space division (separation) switching network — a switching network in which the transmission paths are physically distinct.

Crosspoint — a two-state switching device, possessing a low transmission impedance in one state and a very high one in the other.

Switch — a rectangular array of crosspoints in which one side of the crosspoints is multiplied in rows and the other in columns.

Stage — those switches in a switching network which have identical, parallel functions.

Grid — a two-stage switching network in which a single path exists between every first-stage switch to every second-stage switch. The number of outputs in each first stage must equal the number of second-stage switches; the number of inputs in each second-stage switch must equal the number of first-stage switches.

Link — the connection between terminals on one switch and terminals on a switch in the next stage corresponding to a single transmission path.

Junctor — any link between central stages of the network.

Concentration — the function usually associated with the first stages of a switching network and characterized by configurations possessing fewer output than input terminals; provided to improve network efficiency when the input terminals carry a light traffic load.

Expansion — the inverse of concentration.

Access — a term indicative of existence of paths within a network configuration from an input terminal to a set of output terminals in absence of traffic; partial access refers to the ability of reaching only a fraction of the output terminals; full access permits reaching all terminals by unique paths; multiple access allows reaching all output terminals in more than one way.

Graph — a graphical representation of all possible paths between two network terminals.

Blocking — inability to interconnect two idle network terminals because some of the applicable links are used for other connections.

Erlang — the traffic unit corresponding to an average of one call present on a traffic carrying facility.

Occupancy — average proportion of time that a traffic carrying facility is busy.

REFERENCES

1. Ferguson, J. G., Grutzner, W. E., Koehler, D. C., Skinner, R. S., Skubiak, M. T., and Wetherell, D. H., No. 1 ESS Apparatus and Equipment, B.S.T.J., this issue, p. 2355.
2. Danielson, D., Dunlap, K. S., and Hofmann, H. R., No. 1 ESS Switching Frames and Circuits, B.S.T.J., this issue, p. 2221.
3. Feiner, A., Lovell, C. A., Lowry, T. N., and Ridinger, P. G., The Ferreed — A New Switching Device, B.S.T.J., **39,** Jan., 1960, p. 1.
4. Feiner, A., The Ferreed, B.S.T.J., **43,** Jan., 1964, p. 1.
5. Joel, A. E., Jr., An Experimental Switching System Using New Electronic Techniques, B.S.T.J., **37,** Sept., 1958, p. 1091.
6. Descloux, A., *Delay Tables for Finite- and Infinite-Source Systems*, McGraw-Hill, New York, 1962.
7. Grantges, R. F., and Sinowitz, N. R., NEASIM: A General-Purpose Computer Simulation Program for Load-Loss Analysis of Multistage Central Office Switching Networks, B.S.T.J., **43,** May, 1964, p. 965.