THE BELL SYSTEM TECHNICAL JOURNAL

DEVOTED TO THE SCIENTIFIC AND ENGINEERING ASPECTS OF ELECTRICAL COMMUNICATION

Volume 54 | May–June 1975 | Number 5

Digital Data System:

System Overview

By N. E. SNOW and N. KNAPP, JR.

(Manuscript received July 12, 1974)

This paper presents an overview of the Digital Data System. The services to be provided are described and the premises for establishing both service and system objectives are discussed. The network concept, planning for its growth and administration, and a description of the network elements are presented as an introduction to the detailed papers which follow.

I. INTRODUCTION

The Digital Data System (DDS) is a new data communications network that is integrated into the nationwide telecommunications system. With this network, *Dataphone*® digital services are available. Point-to-point and multipoint private line services are provided. The first allows digital communication between two subscriber terminals, while the latter allows several terminals at different locations to share a common transmission channel. Data rates of 2.4, 4.8, 9.6, and 56.0 kb/s are offered. These services are similar to those which have been available for the past decade using analog telephone channels of various bandwidths. The new system, however, utilizes station-to-station digital transmission techniques as contrasted to modulation and demodulation of digital signals to and from analog form for transmission over telephone channels.

Existing business machine terminals are directly usable on the system through industry standard interfaces or, if the customers prefer, connection may be made at a four-wire channel service interface. Data transmission is synchronous; timing signals are always supplied from the network.

The DDS has become practical and desirable through the large-scale deployment of digital transmission systems in metropolitan areas, the development of long-haul digital transmission systems, progress in technology, and the development of a sufficient market for the services. It offers economies in the cost of transmission and, by taking advantage of the regeneration, monitoring, and protection approaches applicable to digital signals, it promises a higher-quality service than has been realized in using analog telephone systems for data communication.

The system is complex in that it comprises many elements of hardware, planning, operation, and administration. A unique set of abbreviations and acronyms has naturally developed. The appendix is a glossary of those used in this and the accompanying papers.

II. A DDS SERVICE

A point-to-point *Dataphone*® digital service in which a customer's channel interconnects two stations located in different cities is shown in Fig. 1. Three main parts of the DDS are readily identifiable: a local distribution system that makes use of readily available telephone distribution cable pairs to reach the subscriber, a metropolitan area network of digital lines for collecting customer channels from many serving central offices into one "hub" office which acts as the serving test and administration center, and the intercity network of long-haul digital transmission facilities.

From a station, the customer's channel is carried to a local serving office by a newly designed loop transmission system which includes station equipment, a four-wire loop (cable pairs), and an office channel unit which processes the data signals into a format for entry into the network. The local serving office serves as a collection point for individual channels.

The channel is then combined with other channels terminating in the same serving office by a time-division multiplexer into one high-speed signal for efficient transmission over a digital line to a hub office. A hub office acts as a collection point for channels coming from numerous serving offices in the metropolitan area. In the hub office, the channel is separated from other channels by demultiplexing it from the incoming high-speed bit stream to provide maintenance test access and to allow it to be connected to the desired long-haul transmission system.

The channel is then combined with other channels destined for the same distant city by a time-division multiplexer into a single high-

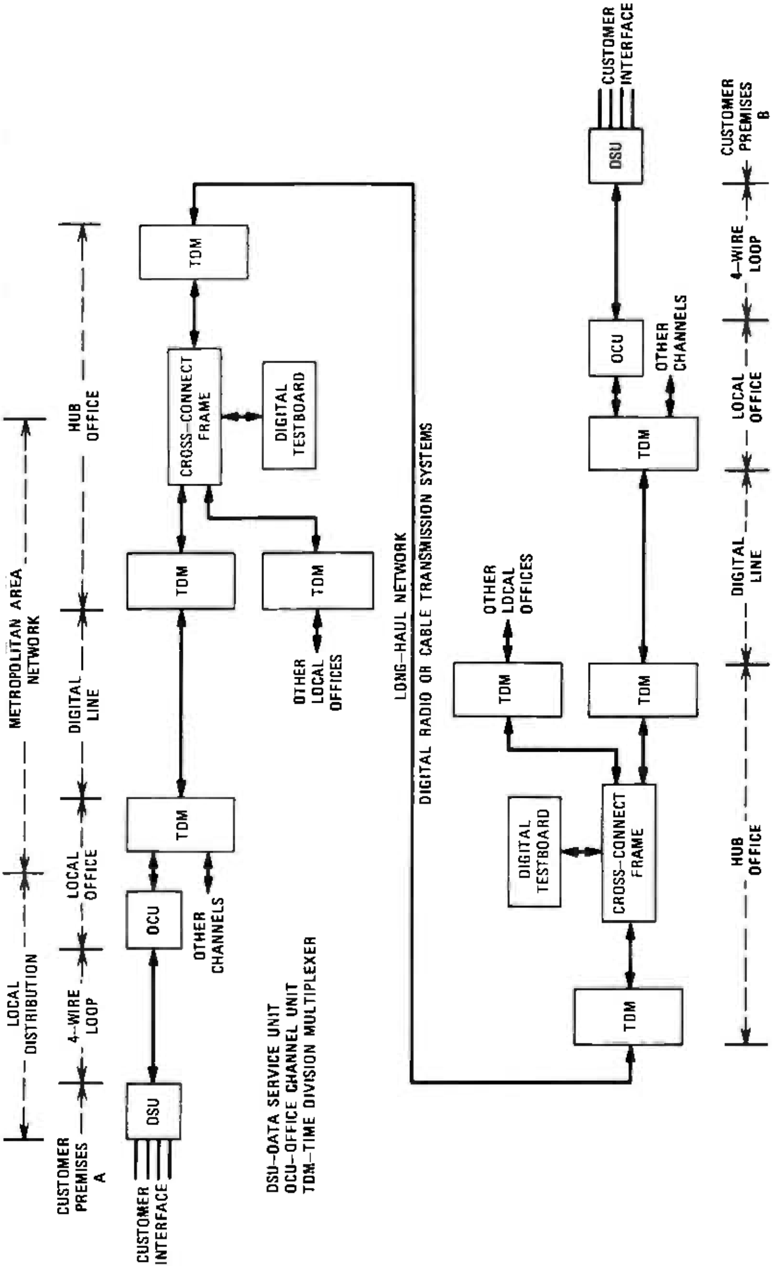

Fig. 1—A point-to-point DDS service.

speed signal for long-haul transmission through the intercity network to the distant hub office. The high-speed signal may be carried over radio or cable systems of several types.

When the high-speed signal reaches the hub office in the destination city, the individual channels are again separated and each is made available for maintenance test access. The customer's channel is then routed to the desired serving office via multiplexers and digital lines in the metropolitan area network and finally to the second station by means of a local loop-transmission system.

This example shows the most simple point-to-point channel including a single link in the long-haul intercity network between hub offices and one local office in each of the two cities involved. In reality, a customer's channel may pass through several local and hub offices of different kinds while traversing the network and may be demultiplexed and remultiplexed several times to efficiently load each link encountered.

III. NETWORK CONCEPT

It is planned that the DDS will rapidly grow into a nationwide network arranged in a three-level facility hierarchy. The highest-level regional hub offices are interconnected by large cross-section transmission facilities and are located along major existing radio and cable transmission routes. Second-level sectional hub offices home on and have transmission links to only one regional hub. Third-level metro hub offices, in turn, home on only one sectional hub office. Within this three-level framework, the geographical area surrounding a hub office is designated either a class I, class II, or class III digital serving area (DSA) corresponding, respectively, to the level of hub office. Facility engineering rules have been developed to provide the required digital transmission capacity between any DSAs in the network in a manner that realizes efficient loading of the facilities.

Each hub office contains time-division digital multiplexers, timing supplies, digital testboards, cross-connect arrangements for flexible channel interconnection, and customer loop terminations for the immediate geographical area. Transmission throughout the network is synchronous. One regional hub office will ultimately contain the timing supply that acts as the master source for the entire network. This supply will be locked to the Bell System reference frequency standard. Timing is derived in each office from selected incoming communication bit streams and is successively passed to equal or lower-level offices in the same manner. This approach synchronizes the entire network by creating a tree-structured timing network.

Homing on the hub offices of the three-level intercity network are the metropolitan networks of serving central offices. An example is shown in Fig. 2. Local serving offices are designated either as intermediate offices or end offices. They are at the end branches of the tree-structured timing network and derive timing information, in the same manner as higher-level hub offices, from selected incoming bit streams. The manner in which local offices are engineered and designated is determined by the characteristics of a particular metropolitan area and the network configuration best suited to the existing central office locations and subscriber distribution. As in the hub offices, local offices contain digital multiplexers, timing supplies, and subscriber loop terminations.

Standard telephone loop plant is used from the nearest DDS central office, local or hub, to reach the subscriber premises. Station terminal equipment is provided to operate at the service data rate and with the appropriate interface.

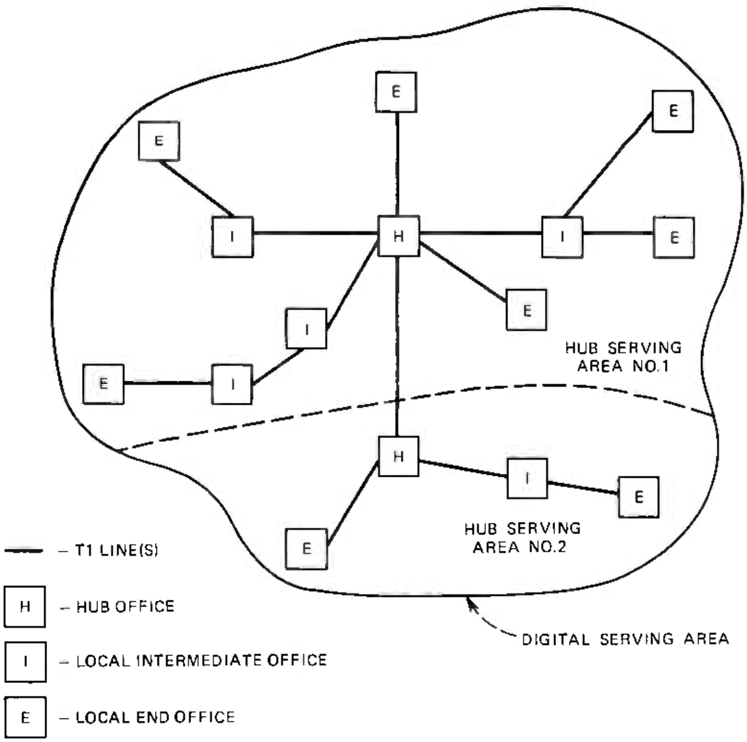

Fig. 2—Example of DDS metropolitan network.

All long-haul facilities in the network are protected by the normal radio or cable protection-switching systems. Digital line facilities in the metropolitan area plant are continuously monitored and protected by standby lines. Time-division multiplexer terminals are continuously monitored and, if a failure occurs, protection equipment automatically takes over. Service-affecting failures result in major alarms, while failures covered by protection equipment result in minor alarms calling for maintenance effort.

Transmission testing for installation or maintenance of subscriber channels is carried out primarily from testboard positions located in hub offices where each channel may be accessed with digital test equipment. Remotely controlled loop-around features activated by digital control code signals make it possible for a testboard operator to isolate trouble to the station, loop, or network without aid from either the subscriber or another testboard operator. Most troubles in or between offices are detected and appropriate personnel informed by alarms before the customer has reported a trouble.

Facility assignment, monitoring, and restoration activities are carried out by the same centralized approaches as have been recently evolving for voice telephone facilities. Within metropolitan areas, this is primarily through T-carrier restoration control centers, and, in the long-haul plant, by regional operating control centers.

IV. OBJECTIVES

New approaches have been taken in establishing both service and system objectives for the DDS.[1] It is recognized that the long-term average bit error rate is not a complete characterization of the performance of a data communication service. For the user to plan effective communication systems, a more useful characterization of performance includes (*i*) the expected amount of the time the service will be available for use, and (*ii*) knowledge of how error events are distributed in time while the service is available. These are the terms in which the objectives for the DDS have been established. Equipment designs, maintenance approaches, and administrative procedures have been directed toward achieving the objectives.

In establishing error performance objectives, the causes of errors were considered. These include such events as protection switches of terminals or transmission links and radio transmission fades. From past experience, it is known that errors usually occur in bursts that can cause considerable variation in the measured average bit error rate but have little effect on the efficiency of the communication. An objective has been established to provide transmission which is error-free in 99.5 percent of all one-second intervals. This objective now offers guidance, for example, in the design of terminals that use block retransmission for error correction and places a bound on the overall throughput efficiency that can be expected.

In establishing service-availability objectives, the kinds of failures that can occur as well as experience in restoring such failures have been considered. An objective has been established that a customer's service will be available for his use an average of 99.96 percent of the time. Allowable outage time is allotted to the various parts of a customer channel. After realistic experience values of repair or restoration time are applied to the transmission systems, subscribers' loops, etc., the remaining time within the objective is allotted to new parts of the system. Where necessary, automatic protection features have been included in terminal or transmission system designs to eliminate most causes of system outage. Maintenance alarm and testing features with appropriate procedures are designed into the system for isolation of trouble conditions and initiation of repair based upon allocations within the objectives.

Planning the system design and operation, while based upon new approaches to establishing objectives, has followed the concepts of centralized administration and restoration control that have been introduced into the telephone network. DDS signal transmission formats are identical to those of PCM voice systems, so they require no special treatment or recognition for either maintenance or administrative activities.

V. NETWORK PLANNING

Communities of interest requiring data communication services are geographically widespread. For the DDS to provide effective communication systems, the network must grow to reach almost every part of the country. Growth must be simultaneous in two dimensions: the interconnection of many metropolitan areas by long-haul digital routes and the penetration of each area to a large number of serving central offices.

Developing a viable growth plan[2] requires the determination of not only a facility network but also the rate of growth in each dimension if the available resources are to be properly used. Estimates of the potential market and location, and information on available or planned transmission routes, have been the basic inputs to the planning studies. Computer aids have been developed to assist in optimizing the initial and growth configurations of the intercity and metropolitan area portions of the network. The output of these aids is directly usable in engineering the system, in forecasting, and ultimately in programming equipment manufacture.

Many objectives and features of the network have led to new areas of planning. For example, maintaining the network synchronization plan requires long-term centralized control of the plan and rapid availability of information to permit connection of new offices or routine rearrangements. New forms of circuit-layout cards specifically designed for digital channels must be available quickly at numerous locations so that tests may be conducted and trouble isolated within objective time limits. Time-shared computer approaches are being utilized where possible for these and other areas of network planning and administration.

VI. NETWORK ELEMENTS

Development of the DDS network concept required the formulation of a plan to use the digital transmission systems already available, or planned for various other services, and definition of the required equipment or systems specific to data services and their unique objectives. The T1 carrier DS-1 digital signal rate (1.544 Mb/s) offered a basic digital capacity available in both short-haul and long-haul transmission plant. Subscriber loop cable plant offered the practical way of reaching the subscriber premises. The additional systems requiring design and development effort to implement the network were the loop and station transmission system, time-division multiplexers and protection equipment, channel interconnection arrangements, test access and test equipment, and timing supplies for signal synchronization.

Figure 3 shows the relationship of network elements below the 1.544-Mb/s (DS-1) digital transmission level.

6.1 Loop transmission system

The subscriber loop transmission system[3] utilizes twisted-pair cable facilities to provide full-duplex four-wire transmission between the subscriber's premises and his serving DDS office. The loop is terminated in either a data service unit (DSU) or a channel service unit (CSU) at the subscriber's location and in an office channel unit (OCU) at the central office. A DSU interfaces a business machine terminal with industry standard control, timing, and data leads and performs all the signal shaping, encoding, and decoding necessary to communicate with the network. A CSU provides only the necessary circuitry to properly terminate the loop with a well-defined interface and allow it to be tested from a central office location. Where a CSU unit is provided, the timing recovery and signal encoding and decoding is incorporated into the business machine terminal. Transmission on the loop is bi-

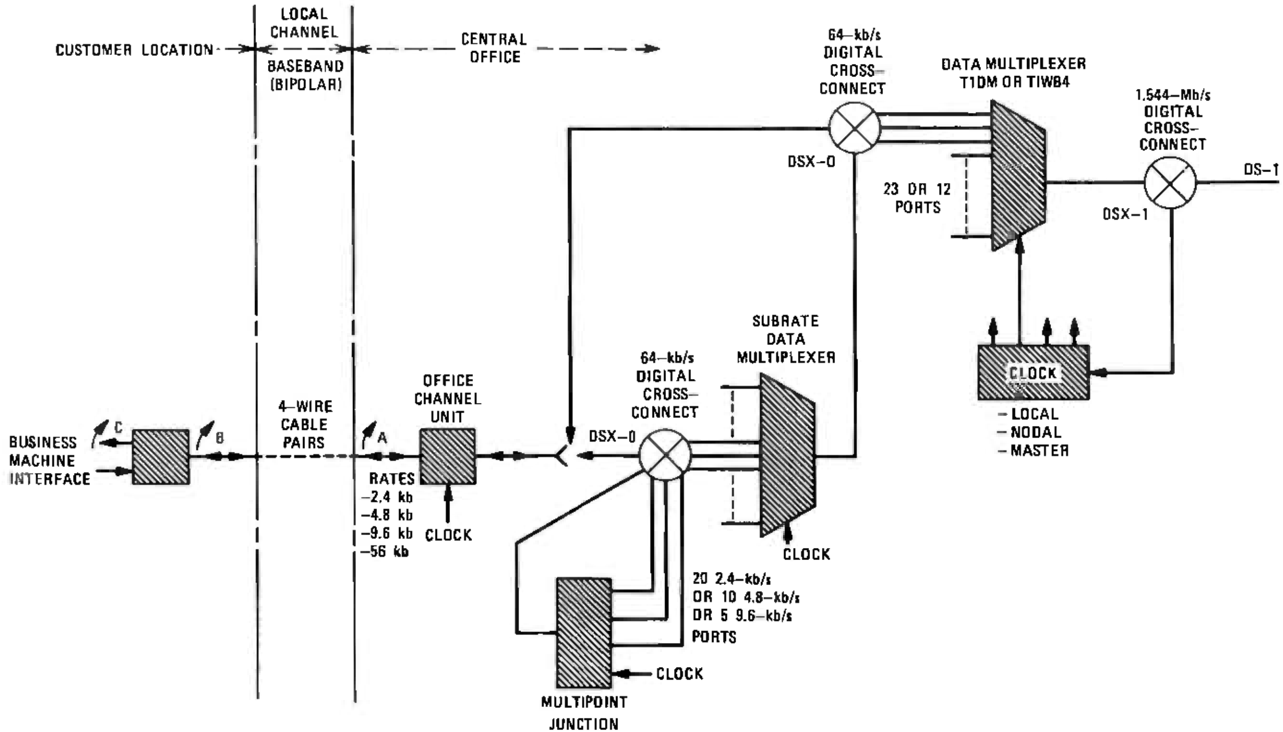

Fig. 3—DDS network elements below DS-1 level.

polar baseband with specific violation patterns to encode test and supervisory control information into the bit stream. Pulse rates are at the service data rate. Transmitting power levels are chosen to minimize the problems of pair selection and coordination with other services in the same cable sheath. The OCU performs similar functions to those of the DSU toward its loop side. On the central office side, it forms the data signals into a "byte" format, adds the necessary control information, and, regardless of the service data rate, builds the signal up to a universal 64-kb/s interconnection rate that has been designated as a "DS-0" signal. This prepares the signal for interconnection to multiplex ports, junction units, or other subscriber loop terminations.

6.2 Time-division multiplexers

A two-stage synchronous multiplexer organization has been developed.[4] The first stage is available in arrangements that accept either twenty 2.4-, ten 4.8-, or five 9.6-kb/s service rate channels, each having been converted to the DS-0 format by the office channel unit, and deliver a single DS-0 signal. Since the multiplexing organization is synchronous, the required rate-changing is a simple process of "byte" repetition or deletion. Two types of second-stage multiplexers are provided: the T1DM which combines up to 23 data channels into a DS-1 signal format, and the T1WB4 which, in conjunction with a D-type channel bank, combines a flexible mixture of up to 24 data and voice channels into the same format. The latter type provides less equipment redundancy and is more economically suited to central offices serving a small number of DDS subscribers. The channels derived by either will accept the DS-0 signal from office channel units, first-stage multiplexers, a 56-kb/s service rate channel, or other multiplexer ports. Modular hardware design of the multiplexer terminals allows complete freedom to equip only the channel capacity required, and to change capacity by addition or removal of plug-in circuit modules without service interruption. Performance-monitor equipment continuously scans the multiplexers in an office and, if a persistent failure is detected, causes protection circuitry to be switched into operation.

By application of the two stages of multiplexing, a flexible mix of customer channels may be derived from a 1.544-Mb/s DS-1 signal which ranges from 23 operating at 56.0 kb/s to 460 operating at 2.4 kb/s. The economics are evident when a comparison is made to the 24 voice-frequency analog channels normally derived by a D-type channel bank from a DS-1 signal, which could be used as 24 data channels operating at 2.4 kb/s.

6.3 Channel interconnection

Interconnection between loops or channels in an office is always accomplished at the 64-kb/s (DS-0) signal rate. In local offices, this may be through jack and connector panels, while in hub offices the interconnection is by plug-ended jumpers on a DSX-0 cross-connect frame. The former provides test access and limited cross-connect flexibility, while the latter provides total flexibility of equipment assignment, channel interconnection, and introduction of jack field appearances for test access.

Since all signals at the DSX-0 cross-connect are in the universal 64-kb/s format, any channel may be cross-connected to any multiplex port of equal or higher rate designation. For example, a 4.8-kb/s channel can be connected to a 4.8-, 9.6-, or 56-kb/s port, but not to a 2.4-kb/s port. This feature makes it possible to minimize or eliminate the installation of first-stage multiplexers until the number of channels required is large enough to necessitate efficient use of the line capacity.

6.4 Multipoint channel arrangements

Multipoint junction units, interconnected at the DSX-0 cross-connect, permit a number of channels and/or loops to be associated with one multipoint communication system. This arrangement is similar to a full-duplex telegraph hub in concept. The maintenance testing features, however, allow remote selection of "legs," which facilitate one-man testing of a multipoint service.

6.5 Synchronization timing supplies

All multiplexers, channel units, junction units, and test equipment within an office operate synchronously from one timing supply which derives its frequency information from a selected incoming DS-1 facility.[5]

The timing supply in a hub office is designated a "nodal timing supply." The frequency of this supply is inherently contained in transmitted DS-1 signals, and is therefore passed on to other offices of equal or lower level in the transmission hierarchy. The hardware is totally redundant, since failure would disable all communication through the office. Occupying so strategic a position in the network demands that the nodal timing supplies employ highly stable oscillators with memory so that loss of incoming frequency information does not disrupt or degrade performance on other facilities and channels through the office before restoration or repair can be effected.

While the network is growing geographically, a nodal timing supply will act as the master timing supply. Its location in the network will change from time to time. Ultimately, one such supply will be connected to the Bell System reference frequency standard and become the "master timing supply."

The timing supply in local offices is designated a "local timing supply." It is, in principle, identical to a nodal timing supply with the exception that lower-cost, less stable oscillators are used. Redundancy is again employed to insure reliability and, since local offices are likely to be unattended for relatively large periods of time, automatic switching to a secondary source of incoming frequency information is provided if such a source exists. A local timing supply may pass frequency information to other local offices, but never back to a higher-level office in the network.

6.6 Maintenance testing

A new digital testboard that will be located in facility hub offices is provided for installation and maintenance testing of DDS channels.[6] A 450-channel-capacity testboard position provides a six-jack (four transmission and two monitor) appearance for each channel. The full duplex channels may be monitored in either direction of transmission, or the channel may be opened and test signals transmitted and received in either direction. Remote loopback tests may be conducted around the loops shown as A, B, or C in Fig. 3. The loopback connections are activated and maintained by control signals and test data generated by a digital transmitter test set. Return signals are observed by a digital receiver test set. The test sets are an integral part of the testboard design and are also available as portable test sets.

The testboard has a number of additional features suited to its serving test center (STC) functions. These include the ability to generate unique test codes which may be connected to a channel for extended periods of time when required in tracing a channel through the network to locate a problem condition. A multipoint signaling unit (MSU) provides a method of selecting a route through remote multipoint junction unit (MJU) branches until a particular station is reached. Loopback tests may then be conducted as with one end of a point-to-point channel. Responses from each MJU encountered along the way verify its identity and the selected branch by means of a numeric display on the MSU. A telephone circuit with key equipment is provided to pick up a number of order wire or dial lines for voice communication. All connections of test equipment to channel appearances are made by means of retracting cord reel circuits. Display of status and test results is by a combination of light-emitting-diode lamps and light-emitting-diode numeric displays.

Use is made of the portable test set versions of the digital transmitter and digital receivers in office equipment areas away from the testboard. Access is provided by test points on channel, junction, and multiplexer units where individual customer channels or selected channels from a multiplexed bit stream may be observed. In local offices, jack and connector panels in the equipment bays provide access to individual channels where loopback tests may be conducted in the same manner as from a testboard.

The one-man testing capabilities designed into the system will enable craftspeople to quickly isolate trouble in a customer's channel and will, in conjunction with the many equipment alarms, aid in identifying the type of maintenance required.

6.7 Facilities

Figure 4 depicts the digital transmission facilities that will be used in the DDS network. The hierarchy of facilities begins with the DS-1 (1.544-Mb/s) signal rate of the T1 carrier line. Within metropolitan areas, the T1 line will actually be the primary facility used to interconnect local end offices and local intermediate offices to the hub office.

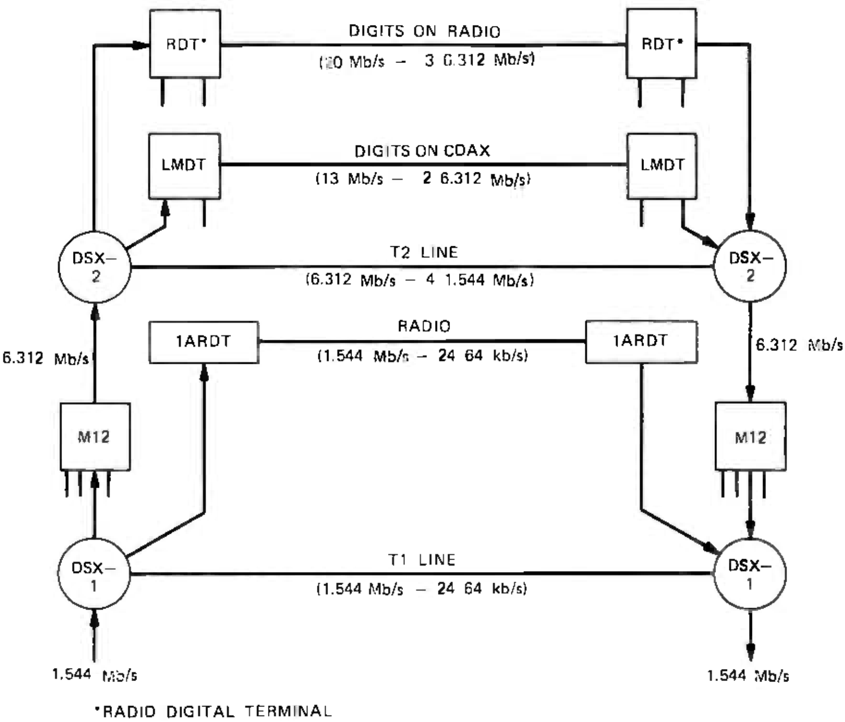

Fig. 4—Digital transmission hierarchy.

Hub offices at all levels will ultimately be interconnected with a variety of long-haul digital systems. During the early years of network growth, the system most used will be the 1A radio digital system (1ARDS) which operates at the DS-1 signal rate. This system, also referred to as DUV (data under voice), employs multilevel signal encoding and shaping to compress the baseband spectrum to a bandwidth that occupies the bottom 500 kHz of the radio baseband. In this manner, a 1.544-Mb/s digital capacity may be derived from each radio system in a route without reducing the message channel capacity of the system.

As the network route cross-section requirements grow beyond the capacity available by application of 1ARDS, it will be necessary to utilize higher-capacity systems. Several have been developed, although not yet deployed. Up to four DS-1 signals may be combined by the M12 multiplexer to form a DS-2 signal of approximately 6.3 Mb/s. In some cable routes, this signal will be directly applicable to the T2 digital line. Further combination of two DS-2 signals (comprising eight DS-1 signals) by the M2L multiplexer produces a signal that will, by use of the L-mastergroup digital terminal (LMDT), occupy one mastergroup band of either the L4 or L5 coaxial cable carrier system.

Another alternative that has been successfully demonstrated for large cross-section digital capacity occupies a full radio channel and has a capacity of approximately 20 Mb/s. By use of M12 multiplexers, this system can provide up to 12 DS-1 facilities.

The range in digital capacity of the available transmission systems and those currently being developed will allow selection to fit the needs of the DDS network during its early years of growth and to meet the needs for the future.

All multiplexing above the DS-1 signal level is asynchronous, using bit-stuffing techniques for rate synchronization. The DDS utilizes hub-to-hub synchronization at the DS-1 rate; therefore, it is not necessary to tie the DDS synchronization network to any of the higher-level multiplexers or transmission systems.

VII. PHYSICAL DESIGN

A key consideration in the physical design[7] of DDS central office equipment is flexibility: (*i*) to permit engineering, ordering, and installation to meet service needs, (*ii*) to grow easily with short installation intervals, and (*iii*) to allow rearrangements for changing service needs with little or no downtime.

Initial and growth installations are greatly simplified by the use of factory-supplied connectorized cables to interconnect bays of equipment or assemblies within a bay. Installation time and chance of installer wiring error are greatly reduced.

Considerable flexibility has been realized within unit designs and the total number of circuit pack codes held as low as possible by devoting special attention to the physical partitioning of circuit functions. For example, the office channel unit (OCU) assembly can function at any of the four data rates or mixes thereof by inserting appropriate plug-in circuit packs. Similarily, the subrate data multiplexer (SRDM) assembly may be equipped with a flexible mix of 5, 10, or 20 port multiplexers for 9.6, 4.8, or 2.4 kb/s channels by selecting the proper circuit packs and the connectors into which they are plugged. Where different channel operating rates are involved within the same equipment assembly, color coding of circuit-pack face-plate labels is used for identification.

Care has been taken to ensure reliability of channel connections at the DSX-0 cross-connect where the interconnection of 64-kb/s signals among channel units, multiplex ports, etc., is accomplished as a continuing activity. The cross-connect panels mount on duct-type bays that can be located in the same lineup with other equipment bays. Within the panels, four-wire plug-ended jumpers engage four recessed pins in a plastic cell structure. Seated jumper plugs are automatically locked in place to avoid accidental disconnections. Removal of a jumper can be accomplished only with a special tool that is secured to the bay.

Precautions have been exercised in developing cabling, shielding, and grounding arrangements to ensure performance in the wide range of operating environments to be encountered. Designs have been realized that make it easy to introduce the required service capacity into existing telephone central offices. While fixed equipment configurations have been coded for convenience in engineering and ordering, the connectorized assemblies may be mounted on a miscellaneous basis if special arrangements are desired by an operating company.

VIII. SUMMARY

The services, the objectives, and the elements for realization of the Digital Data System have been outlined. The papers that follow describe in more depth the systems that have been developed, the strategies for their use, and the planning being carried out to provide the features of a specialized digital data communications network, while utilizing to a maximum the existing Bell System network and its established operating methods.

IX. ACKNOWLEDGMENTS

Many individuals have contributed to the planning and realization of the Digital Data System. Special recognition is due T. H. Thompson and R. L. Wagner who in the course of the development program provided the project organization and leadership, and to C. R. Moster, L. R. Pamm, J. T. Bangert, and U. S. Berger in whose laboratories or engineering centers the major planning and system development efforts were pursued.

APPENDIX

Glossary of Acronyms and Abbreviations

Throughout this series of papers, certain terms, abbreviations, and acronyms are used to simplify and shorten the presentation of ideas. A listing of some more commonly used terms relating to the Digital Data System is presented here with brief descriptions.

ALBO Automatic line-build-out network.

AVAILABILITY Percentage of time that satisfactory data communication service is available. The term "satisfactory" implies that terminal equipment and cables are in working order.

BASEBAND In the Digital Data System, a digital stream designated to contain data for only one customer station. For example, data on a customer's loop or at the DS-0A level are at baseband, while data at the DS-0B level are not at baseband.

BASEBAND OFFICE An office in a DDS digital serving area that contains no DDS multiplexing gear, but acts as a link-up point between the four-wire connection to a customer's station and an interoffice cable.

BCPA (bay clock, power, and alarms shelf) A DDS equipment shelf used in conjunction with office timing supplies to supply timing to equipment bays. It also supplies power to equipment in the bays and combines alarms.

BIPOLAR RZ (BPRZ) (bipolar return to zero) A three-level code in which alternate 1s change in sign (for example, 1011 becomes +1, 0, −1, +1) and transitions between adjacent 1s pause at the zero voltage level.

BIPOLAR NRZ (BPNRZ) (bipolar nonreturn to zero) The same as bipolar RZ, except that transitions between adjacent 1s do not stop at zero level.

BPV (bipolar violation) A violation of the alternating +1, −1 pattern in a three-level code.

B6ZS (bipolar with 6 zero substitution) A coding scheme, implemented by the M12 multiplex, whereby any group of six consecutive zeros is converted into a known bipolar violation pattern at the DS-2 level.

BYPASS CIRCUIT A DDS circuit that is routed directly from DSX-0B in the local access multiplexing section of a hub office to DSX-0B in the long-haul access multiplexing section without appearing at DSX-0A.

BYTE In the Digital Data System, a group of eight consecutive binary digits associated with a single user.

BYTE STUFFING In the Digital Data System, the technique by which the bit rate of a digital stream is increased by repeating bytes and transmitting them at a faster rate. The information content of the stream is not increased.

CHAIN OFFICE A local end office having a T1WB4 at one of the intermediate points connecting two links in a T1WB4 chain.

CHAIN See *T1WB4 chain.*

CONTROL SIGNALS Signals in byte format used for synchronization, status, and remote testing.

CROSS-CONNECT A piece of hardware used to interconnect multiplexers with line-terminating equipment and other multiplexers. Access to signals is often available through jacks associated with a testboard located near the cross-connect.

CP (circuit pack) A unit that contains part of the DDS circuitry and can be inserted into equipment shelves where required.

CSU (channel service unit) A unit located on the customer's premises that terminates a DDS channel and is used with the customer's logic and timing recovery circuitry.

DATA MODE A condition of the DSU with respect to the transmitter in which its data-set-ready and request-to-send circuits are ON and it is presumably sending data.

DDS Digital Data System.

DDS LOOP That portion of an individual customer's channel between the station and its associated office channel unit (OCU).

DMC (data-message combiner) Combines analog multiplexed voice signals with a DS-1 level data signal.

DOWN TIME Time during which data communication is not available or is unsatisfactory (see *Availability*) because of malfunction. Time required for preventive maintenance is not included.

DSA (digital serving area) The combined geographical serving areas of a set of DDS serving offices, as specified in the appropriate tariff(s). The DDS office serving areas making up a DSA are not necessarily contiguous, and a DSA may overlap state and associated company boundaries; however, a typical DSA might encompass one urban area of a single associated company.

DSU (Data Service Unit) A terminal located on the customer's premises for the purpose of accessing the Digital Data System through a standard Electronic Industries Association (EIA) or Comité Consultatif International Télégraphique et Téléphonique (CCITT) interface.

DS-0 (digital signal at the 0th level of the DDS TDM hierarchy, the DS-0 level) A signal at the 64-kb/s rate (the DS-0 rate).

DS-0A A DS-0 signal designated to carry data for only one station. For subrate speeds, successive bytes are repeated as necessary to match the customer's data speed. Only DS-0A data signals appear at the DSX-0A cross-connect.

DS-1 (digital signal at the first level of the DDS TDM hierarchy, the DS-1 level) A biopolar return-to-zero signal at a 1.544-Mb/s rate (the DS-1 rate).

DS-2 (digital signal at the second level of the DDS TDM hierarchy, the DS-2 level) A bipolar return-to-zero signal at a 6.312-Mb/s rate (the DS-2 rate).

DSX-0 (hub X-conn) Digital cross-connect used to interconnect equipment at the DS-0 level. Note that no cross-connects are used in DDS local offices.

DSX-0A (STC X-conn) The DSX-0 digital cross-connect at a DDS hub office where individual customer circuits are properly routed.

DSX-0B (multiplex X-conn) The DSX-0 digital cross-connect at a DDS hub office used to connect T1DM and T1WB4 ports with SRDMs and to connect T1DM and/or T1WB4 ports together for through or bypass circuits.

DSX-1,2,3 Digital cross-connect used to interconnect equipment, provide patch capability, and provide test access at the DS-1, DS-2, or DS-3 level respectively.

DUPLEX A communication mode in which transmission can occur in both directions simultaneously (sometimes referred to as full duplex).

DUTY CYCLE The percent of a single pulse period (for a 1) during which the voltage is nonzero.

DUV See *IARDS*.

EFFICIENCY OF DATA COMMUNICATIONS Percentage of one-second intervals in which data are delivered free of error.

EFS Error-free seconds.

END OFFICE In a digital serving area, a local office that passes on toward the hub-only circuits that entered the office over local loops. The main function of an end office is to combine several individual customer channels, by means of DDS multiplexers, and to transmit the combined bit stream toward the STC hub.

FMT/FMR (FM transmitter and FM receiver) Used with broadband radio systems.

FOUR-WIRE CIRCUIT A facility that provides two full-time, independent channels for transmission in opposite directions. It is historically associated with two wires for transmission and two wires for reception.

FRAME On a T1 line, 193 binary dibits, that is, 24 bytes plus one framing bit.

HALF DUPLEX A facility that permits transmission in both directions, but only one direction at a time.

HIT Any disruption of service that persists for less than one second.

HUB An office in the Digital Data System that combines the DS-1 data streams from a number of local offices into signals suitable for transmission over DDS facilities at the DS-1 level or above. Cross-connects at a hub are made via DSX-0. Also see *STC hub*, *Collection Hub*, *Regional Hub*, *Sectional Hub*, and *Metro Hub*.

HUB CROSS-CONNECT See *DSX-0*.

IDLE CODE A bipolar violation sequence transmitted by the DSU to indicate no data are being sent over the line.

IDLE MODE A condition of the DSU with respect to the transmitter in which its data-set-ready circuit is ON, but its request-to-send circuit is OFF, and it is sending idle code.

INTERMEDIATE OFFICE In a DDS digital serving area, a local office that passes on toward the hub circuits that entered the office over a T1 line, in addition to those that entered over local loops.

ISMX (integral subrate multiplexer) A subrate multiplexer arrangement used only in local offices, in which the subrate multiplexing function is contained within the OCU shelves.

JCP (jack and connector panel) A unit used in a local office to connect the various equipment pieces and to provide test access with portable test sets. See *M-JCP* and *SM-JCP*.

kb/s Kilobits (10^3 bits) per second.

LDFMC (long-distance facilities maintenance center) A long-lines toll test room concerned with maintenance of long-haul digital systems.

LMDS (L-mastergroup digital system) A system that provides for the transmission of two DS-2 signals in one of the mastergroup bands of the L4 or L5 coaxial cable systems.

LMDT L-mastergroup digital terminal.

LOCAL ACCESS MULTIPLEXING The multiplexing equipment in a DDS hub office dedicated to combining circuits for transmission to local offices or another hub office in the same digital serving area.

LOCAL LOOP The cable pairs between a DDS office and customer premises.

LOCAL OFFICE A DDS office that concentrates "on-net" customer circuits into T1 streams which can be transmitted to a hub office. Cross-connects at a local office are made via JCPs.

LTS (local timing supply) Common timing source for a DDS local office. In the absence of input timing information, this unit is less stable than the NTS.

LONG-HAUL Transmission distances typically beyond 50 miles utilizing, for example, TD, T2, L1, or L5 facilities.

LONG-HAUL ACCESS MULTIPLEXING The multiplexing equipment in a DDS hub office dedicated to combining circuits for efficient transmission to *other* local serving areas.

LOOPING (LOOPBACK) A testing procedure that causes a received signal to be transmitted (i.e., returned to the source).

Mb/s Megabits (10^6 bits) per second.

M12 A multiplexer that combines four DS-1 signals into a DS-2 signal.

METRO (CLASS III) HUB A hub office in the lowest of three levels in the interhub routing hierarchy.

MTS (master timing supply) The modified nodal timing supply that receives input timing information from the Bell System reference frequency standard and provides this timing information to the rest of the system.

METROPOLITAN AREA See *Digital Serving Area.*

M-JCP A jack and connector panel that gives access to the ports of a T1DM or a T1WB4.

MJU (multipoint junction unit) A unit employed at a DDS hub office to link together three or more segments of a multipoint circuit.

MULTIPLEX CROSS-CONNECT See *DSX-0B.*

MULTIPOINT A customer circuit with more than two end points. One end point is designated the "control" station.

MSU (multipoint signaling unit) A device used in conjunction with the DDS test equipment to isolate and test various segments of a DDS multipoint circuit.

MUX Multiplexer.

NTS (nodal timing supply) Common timing source for a DDS office. This unit is highly stable in the absence of input timing information, and is only used at hub offices.

OCU (office channel unit) A terminal located in the central office which terminates the customer's loop and provides signal and format conversions between the two types of baseband signals (DS-CS and DS-0A).

OFF-NET A location beyond the primary serving area of the Digital Data System.

ON-NET A location within the primary serving area of the Digital Data System.

1ARDS (1A radio digital system) A system that provides for the transmission of one DS-1 signal over a microwave radio link. This system is also known as data under voice (DUV).

1ARDT (1A radio digital terminal) A digital terminal used in the 1ARDS which converts a T1 line signal to a seven-level partial response format. The resultant signal has a bandwidth of 0 to 500 kHz, and it can be transmitted below the message on a radio facility.

OUTAGE Any disruption of service that persists for more than one second.

PCM (pulse code modulation) The process in which analog signals are sampled, quantized, and coded into a digital bit stream.

PLL (phase locked loop) A circuit containing a variable frequency oscillator whose phase is compared with a reference signal. By a suitable feedback mechanism, both signals are forced to agree in frequency and possibly in phase.

Q (quad) A group of four wires that carry a four-wire circuit.

ROCC (regional operations control center) Coordinates the restoration of failed L4, L5 carrier, or TD/TH radio routes.

SLIP A defect in timing that causes a single bit or a sequence of bits to be omitted or read twice.

SHORT-HAUL Transmission distances typically less than 50 miles.

SM-JCP A jack and connector panel that gives access to the ports of a subrate data multiplexer.

STC SERVING AREA The geographical area for which an STC has maintenance responsibilities.

STC (serving test center) A test location established to control and maintain circuit layout records (CLR), receive customer trouble reports, assist in the checkout of newly installed stations, perform trouble localization, and coordinate service restorals.

STC CROSS-CONNECT See DSX-0A.

STC HUB A hub office that has an STC.

STC HUB SERVING AREA The geographic area covered by all DDS customer stations that home on a single STC hub office.

STRAIGHTAWAY TEST A test procedure in which a test signal is transmitted from one point to a receiver at a different point.

SRDM (substrate data multiplexer) A unit that combines a number of data streams at or below some basic rate (2.4, 4.8, 9.6 kb/s) into a single DS-0B 64-kb/s time-division multiplexed signal.

SUBRATE In the Digital Data System, a data bit rate that is either 2.4, 4.8, or 9.6 kb/s.

TCAC (T-carrier administration center) A center with responsibility for the maintenance and restoration of the T-carrier facilities on an automated basis.

TD-2, -3, etc. A point-to-point microwave radio transmission system.

TDM (time division multiplexing) The process of combining a number of digital signals into a single digital stream by an orderly assignment of time slots.

TEST MODE A condition of the DSU in which its transmitter and receiver are inoperative because of a test in progress on the line.

TRCC (T1-carrier restoration and control center) Performs the same functions as the TCAC, but on a manual basis.

T1 AUTOMATIC STANDBY UNIT (T1ASU) A unit that monitors a regular T1 line and its standby T1 line, and automatically switches to the standby, based on the bipolar violation rate of the regular line.

T1 LINE A digital transmission line that carries data at the 1.544-Mb/s rate (DS-1 level); in DDS, it is used primarily for short-haul links.

T1DM (T1 data multiplexer) A multiplexer that is capable of time-division multiplexing up to twenty-three 64-kb/s channels and synchronizing information into a DS-1 signal.

T1WB4 A voice-data multiplexer capable of combining up to twelve 64-kb/s DS-0B data channels with PCM-encoded voice channels from a D3 or D1D channel bank. The resultant TDM format is a DS-1 signal. Development is under way to permit up to 24 DS-0B data channels, instead of up to 12.

T1WB4 CHAIN An arrangement using T1WB4s to allow a local end office and up to two chain offices to share usage of a single T1 line, which is routed to an STC hub.

T2 LINE A digital transmission line that carries data at the 6.312 mb/s rate (DS-2 level) for distances up to 500 miles.

VOICE MMX (voice mastergroup multiplex) Used to combine 600 voice channels into a spectrum suitable for transmission via broadband radio systems.

X-CONN Cross-connect.

REFERENCES

1. J. J. Mahoney, Jr., J. J. Mansell, and R. C. Matlack, "Digital Data System: User's View of the Network," B.S.T.J., this issue, pp. 833–844.
2. P. F. Brown, Jr., G. W. Phipps, and L. A. Spindel, "Digital Data System: Network Planning," B.S.T.J., this issue, pp. 861–878.
3. E. C. Bender, J. G. Kneuer, and W. J. Lawless, "Digital Data System: Local Distribution System," B.S.T.J., this issue, pp. 919–942.
4. P. Benowitz, S. J. Butterfield, M. P. Cichetti, Jr., and T. G. Cross, "Digital Data System: Digital Multiplexers," B.S.T.J., this issue, pp. 893–918.
5. B. R. Saltzberg and H. M. Zydney, "Digital Data System: Network Synchronization," B.S.T.J., this issue, pp. 879–892.
6. S. M. Fitch and D. L. Rechtenbaugh, "Digital Data System: Testing and Maintenance," B.S.T.J., this issue, pp. 845–860.
7. H. C. Illium, W. B. Lueft, and D. W. Rice, "Digital Data System: Physical Design," B.S.T.J., this issue, pp. 943–964.

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

Digital Data System:

User's View of the Network

By J. J. MAHONEY, JR., J. J. MANSELL, and R. C. MATLACK

(Manuscript received July 12, 1974)

The utilitarian aspect of the Digital Data System is discussed with emphasis on performance objectives that will be important when data communications is inserted into a system of data processing. Objectives for the dependability and quality of data communications are quantified and evaluated in terms of their impact on data processing. Characteristics of the several types of channels available are described in detail along with operational features of particular importance at the interface between data communication and data processing.

I. PERFORMANCE OBJECTIVES

1.1 General

Modern industries are motivated to centralize the control of their integrated operations at a powerful computer and to disperse important operational functions to outlying areas. Modern telecommunications provide the data transmission services that are essential to coordinate dispersed functions with centralized control. Viewed in this light, data communication becomes one link in a larger chain of operations. Objectives for the reliability and quality of data transmission are evaluated in terms of their impact on the larger operation. Data service should generally be deemed excellent when the capability of a data processing system is not diminished perceptibly by the insertion of data communication.

1.2 Availability

The dependability of data service may be expressed in terms of availability, which is simply the complement of average annual down time. The dependability of the service of the DDS is quite naturally related to the inherent reliability of the network of telecommunications facilities from which individual data transmission paths are derived. Dependability of service has been enhanced by providing for automatic substitution of standby facilities during emergencies. The most vulnerable part of the service tends to be the portion of plant that is necessarily dedicated to a single user, where redundant facilities would be most expensive. With automatic substitution to maintain continuity of service during failure of facilities in the common network and with prompt repair of dedicated plant, the design objective for DDS services is to attain average availability of at least 99.96 percent. This objective would permit average annual down time no greater than 210 minutes, which amounts to 0.04 percent of the total time in the 24-hour days of a 365-day year. Thus, communications down time is expected to be a small fraction of the total down time of an entire system of data processing.

The causes of communication failure cover such a wide range of possibilities that it is not possible to specify an amount of down time expected in a particular year for the maximum duration of failure if it should occur.

1.3 Quality

Along with stringent objectives for availability and maintainability, a quantitative description of the quality of data communications is necessary to guide efficient use of the service. It has been traditional to describe the quality of data communication in a simple one-number characterization called error rate, which is the ratio of bit errors to the total number of bits transmitted in a test period. Efficient utilization of data services, however, depends on the exact structure of error patterns as well as the average probability of bit error. For example, where data is blocked into messages containing substantial numbers of bits, the throughput efficiency realized is sensitive to the correlation among bit errors. If bit errors are independent (i.e., no correlation), the number of messages that must be retransmitted to avoid error is very nearly equal to the number of bit errors. Where bit errors are clumped within the block interval, the number of error-free messages is substantially greater than indicated by the average probability of bit error. To minimize the impact of imperfect data communication on the overall system of data processing, the design objective for average efficiency of data communication is to attain 99.5-percent error-free seconds. The percentage of error-free seconds is not related to the more familiar bit-error rate in any simple way, but the proportion of error-free seconds is often estimated by the mathematically convenient relationship,

$$\mathrm{P(EFS)} = (1 - P_e)^B,$$

which gives the probability of an error-free second that would arise from the probability of bit-error P_e at bit rate B. This formula consistently underestimates the percentage of error-free seconds because it relies on complete independence among bit-errors, an assumption known to be invalid. In a practical communications network, transmission is subject to occasional and momentary perturbations, e.g., the automatic substitution of standby facilities to avoid outage of significant duration. Such perturbations seldom persist longer than a small fraction of a second, but they may be interpreted as error bursts. Short-term unavailability of this type is included in allowances of erroneous seconds. Thus, the objective is properly stated in terms of error-free seconds without implication about the ratio of bits found in error to total bits transmitted.

The 1-second sample is short enough to display high resolution in the description of data communication quality, yet long enough to encompass most block lengths that users may choose for error control, and long enough to span most momentary degradations of service. Thus, 1-second samples should approach the behavior of independent events in probabilistic analysis. Since 1 second is regarded as an upper bound on the duration of error bursts, the user may attain transmission efficiency greater than the objective for data communications efficiency if his block length is less than the number of bits transmitted in 1 second. The throughput he actually attains will, of course, depend on irreducible transit time in communication over long distances and on other delays, including those encountered in his own method of error control.

1.4 Allocation of quality

A model of the limiting DDS connection would comprise a long-haul network interconnecting two local serving areas in which there are two local T1 carrier lines in tandem with one baseband loop. The results of field tests, modified by experienced judgment, lead to allocating the total tolerance of errors in the following way:

2 baseband loops would account for 2/10
4 local lines would account for 3/10
1 long-haul network connection would account for 5/10.

The carrier facilities (both long-haul and local) are interfaced at the DS-1 bit rate of 1.544 Mb/s, and the error performance objectives are expressed in error-free seconds at the DS-1 interface for those facilities. The method of allocation is based on the general formula:

$$\text{Percent EFS objective} = 100 - TAP/N,$$

where

T = total tolerance of 0.5 percent, the complement of 99.5-percent EFS

A = fractional allocation

P = number of ports demultiplexed

N = an empirical factor relating the average number of errored seconds expected in DS-0 ports for each one found on the facility.

Appropriate N factors are not yet firmly established but estimates derived from preliminary field tests indicate that 99.6-percent EFS would be applicable to individual local carrier lines. On a comparable basis, the indicated objective would be 99-percent EFS for long-haul network connections. Both objectives would apply at the DS-1 interface where the bit rate is 1.544 Mb/s.

1.5 Performance estimates

Maintainability studies are continuing, but they are necessarily based on hypothetical probabilities of failure and estimates of the time that will be required to restore service. Contributions to annual down time from many of the possible failures have been rendered negligible by providing for automatic substitution of standby facilities. Where automatic protection of service is not provided, system fault-location features have been built in to reduce the duration of service outage by minimizing the time required to locate faults. Study results to date indicate that the objective for availability will be feasible for the large majority of connections expected. Preliminary results of field tests of individual subsystems indicate that the objective for quality of data communications will generally be met.

II. POINT-TO-POINT DDS SERVICE

The duplex private line types of channels are point-to-point and multipoint. As indicated in the article by Snow and Knapp, these are synchronous channels operating at 2.4, 4.8, 9.6, or 56 kb/s. Two choices of terminations are available to the customer. The first is known as the data service unit (DSU) and the second is the channel service unit (CSU) which is part of the basic channel offering.

2.1 Data Service Unit (DSU)

The DSU is physically located on the customer's premises with the output connected to a four-wire loop facility connected to the central office and the input connected to the customer's data-terminal equipment. The DSU consists of two basic sections, a channel terminator, and an encoder-decoder. The function of the channel terminator is (*i*) to provide a balanced termination for the four-wire loops, and (*ii*) to provide the circuitry for implementing the loopback tests. The encoder-decoder consists of a transmitter, a receiver, and a clock recovery section and provides the EIA and CCITT drivers and terminators that interface with the data terminal equipment. The basic function of this unit is the conversion of EIA RS-232-C or CCITT V.35 interface signals to baseband bipolar line signals, and vice versa.

2.2 Data terminal interface

The DSU uses one of two interface connectors, depending on the service offering, one for 2.4-, 4.8-, or 9.6-kb/s service, and the other for 56-kb/s service. For the former, the interface signals exchanged between the data terminal and the DSU are in bipolar voltage form and conform to EIA Standard RS-232-C.

For 56-kb/s service, the clocks and data (transmit and receive) are dc coupled balanced signals. These signals and the interface circuits involved meet the balanced interface standard of CCITT Recommendation V.35. The control signals conform to EIA Standard RS-232-C. For 56-kb/s service, all interface signals should be transmitted over balanced-pair conductors for improved performance and less crosstalk.

For each service offering the interface circuits provided by the DSU match the Type-D interface of RS-232-C for dedicated line service. When the Permanent Request to Send option is used, the DSU has a Type-E interface. Therefore, the DSU provides "plug-for-plug" interchangeability with Type-D or Type-E interfaces of present data sets used on private-line, analog, voiceband networks.

2.3 System operation

The DDS, in the initial service offering, provides for two-point, four-wire, duplex, private-line, digital-data transmission. Although the DDS is a four-wire network, the customer terminals may operate either in a one-way, half-duplex, or in a duplex manner. In describing the operations of the DSU, four modes of operation can be defined: data, idle, out of service, and test. The transmitting section of the DSU can attain the data or idle mode independent of the state of the receiving section, while the receiving section can attain the data, idle, or out of service mode independent of the state of the transmitter section. The test mode involves both the transmitting and receiving section of the DSU.

2.3.1 Half-duplex operation

In half-duplex operation only one terminal transmits at a time. The data terminal desiring to transmit switches its request-to-send circuit on. After a delay, the clear-to-send circuit switches on, indicating that the data terminal may begin transmission. The receiving data terminal has its request-to-send circuit switched off.

To turn the circuit around, the transmitting data-terminal equipment should send an end-of-message (EOM) code and then switch its request-to-send circuit off. Upon receiving the EOM code, the receiving data terminal equipment switches the request-to-send circuit on and after a short delay receives a clear-to-send on signal. If the permanent-request-to-send option is used, the receiving terminal may start transmitting immediately after the EOM code is received. The transmission delay between terminals consists of the propagation delay determined by the routing of the specific circuit and a fixed delay through Bell System terminal equipments. The transmission delay for one-way transmissions over terrestrial facilities will generally be less than 50 ms.

For transmit-only service, it is advisable that the permanent-request-to-send option be used to avoid the clear-to-send delay.

2.3.2 Duplex operation

Since the DDS provides four-wire point-to-point service, and simultaneous transmission in both directions is possible, it is convenient to use the permanent request-to-send option of the DSU so that the clear-to-send circuit is always on. With this option, the data terminal equipment must have a Type-E interface of EIA RS-232-C. When the request-to-send circuit is under the control of the data terminal equipment, the DSU has a Type-D interface.

2.4 Testing and maintenance

The DSU provides testing ability under manual-switch control on the DSU or under the control of the Serving Test Center (STC). When the DSU is in the Test mode, an indication is given to the customer by means of either the LL (local test) lamp or the RT (remote test) lamp.

2.4.1 Manual control of test modes

A test switch provides the customer with the capability of performing an LL or an RT.

2.4.1.1 Local test. With the test switch in the LL position, the DSU is in the local test mode. The LL test permits the customer with a duplex terminal to test the back-to-back performance of his data-terminal equipment and DSU by connecting the transmitter section of the DSU to the receiver section. In addition, the receive line is connected through terminating equipment to the transmit line to allow a signal to be maintained in both directions. For this test the data-set-ready circuit is switched off, but the other control interface circuits, request-to-send, clear-to-send, and received-line-signal-detector, operate as in the idle or data mode.

When the LL test switch is operated, the line is looped in both directions. This gives the remote terminal the capability of testing the transmission path to and from the local DSU as well as permitting the local terminal to test its DSU, as described above. During LL test, the clock of the local DSU is held in synchronization with the system clock.

2.4.1.2 Remote test. With the test switch in the RT position, the DSU is in the remote test mode. In this test mode the output of the received data circuit is connected to the input of the transmitted data circuit at the data terminal interface of the DSU. For this test, the control interface circuit drivers to the data-terminal equipment are switched off and the transmitted data and received data circuits from and to the customer are left open.

With the local DSU in the RT test mode, the remote data terminal has the capability of checking system operation exclusive of the local data terminal. This permits the customer to deduce whether the local data terminal is responsible for a system trouble condition.

2.4.2 Remote control of test modes from the serving test center

In addition to the manual control of the test modes, the telephone company's STC can place the DSU in either the LL or RT mode to test the operations of the line and DSU.

In the RT mode, the STC ascertains whether there are any defects in the transmitter, receiver, and interface circuits of the DSU and the transmission path to and from the customer. It does not ascertain whether the customer is putting proper signals on the interface circuits.

If the results of the RT test show that there is a trouble condition, then the STC can place the DSU in the LL test mode to isolate the trouble condition between the DSU and the transmission path.

2.5 Channel service unit (CSU)

An optional interface to the DSU is available and is known as the CSU. It provides a minimum channel termination that allows for the remote testing of the local DDS channel.

Nominal 50-percent duty cycle, bipolar pulses are accepted from the customer on the data transmit (DT) and data receive (DR) leads. These pulses must be synchronous with the DDS and limited to a specified maximum jitter. The input bipolar pulses are amplified and filtered, and pass through the transmit-repeat coil to the transmit pair. The received signal is first amplified and equalized and then sliced. The resultant bipolar pulses are then passed to the customer. From these pulses, the customer must recover the synchronous clock used for timing the transmit data and sampling the received data.

III. MULTIPOINT DDS SERVICE

Multipoint service has existed in data communications from the earliest days of telegraphy as means for sharing a single channel among several stations which individually do not generate sufficient traffic to justify a full-time dedicated circuit. With multipoint DDS service, three or more customer stations may be connected onto a single circuit that can be operated more efficiently than existing circuits due to lower channel turnaround delays. DDS makes multipoint service available at 2.4 kb/s, 4.8 kb/s, and, for the first time, 9.6 kb/s and 56 kb/s. The customer stations may be located at a number of different customer sites served by DDS and all stations on a single circuit will transmit at the same DDS transmission rate (i.e., 2.4 kb/s, 4.8 kb/s, 9.6 kb/s, or 56 kb/s). All DDS multipoint circuits have two-way simultaneous transmission capability, but can, of course, be used in a half-duplex or one-way-at-a-time fashion. Like two-point DDS circuits, multipoint DDS circuits are transparent (no coding restrictions) to the content and format of the data which the customer transmits.

DDS multipoint service serves only those multipoint circuits that have a single customer-control location and a number of outlying stations. Communication is from any outlying station to the control location and from the control location to any set of outlying stations. Communication from one outlying station to another is not possible except via the control location. All data transmitted by the control location is delivered by DDS to every station on the multipoint circuit.

The customer is responsible for the overall supervision and signaling control of the circuit. His responsibilities include:

(*i*) Inserting addressing information at the control station to permit outlying stations to determine if the information is destined for them.

(*ii*) Detecting the addressing information at outlying stations.

(*iii*) Supervising the communications circuit from the control location to insure that outlying stations do not attempt to transmit simultaneously.

These responsibilities are normally performed by the customer through the use of a communications controller at outlying stations and a computer at the control location.

3.1 Multipoint junction unit

Multipoint service capability in the DDS is provided by interconnecting standard, point-to-point customer channels at the 64-kb/s level by means of Multipoint Junction Units (MJU). To facilitate testing, MJUs are located in DDS hub offices. Although multipoint service is available at all DDS standard data rates, the MJU itself is independent of customer data rates since input and output is at the 64-kb/s level.

MJUs are inserted into a multipoint circuit at hub locations, where they can split the data path from the control station into two or more branches directed toward the outlying stations and combine the data branches from the outlying station into one path toward the control station. The MJU itself is a two-circuit card device that can provide one path toward the control location and two branches on one card, or one path toward the control location and up to four branches on a total of two cards. Additional branches are obtained by cascading MJUs with a cascade of N four-branch MJUs resulting in $3N + 1$ available branches.

Data being transmitted along the communications path from the control location is passed through the MJU unaltered and delivered to all branches toward the outlying stations. Data transmitted from outlying stations to the control location enter a 1-byte serial-shift register whose output is connected to the input of an AND gate. If an incoming byte from a branch is a data byte, it passes unchanged to the AND-gate input. If the byte is a network-control byte, however, it is changed to the all-1's byte when passed to the AND gate. In effect, the control byte is suppressed to a data byte of all 1's. This results in confining possible trouble condition indications to a single branch, thus preventing interference with communication from other branches.

3.2 Circuit operation and implications

There are at least two major ways of operating a DDS multistation line. In the first method, all stations, including the control station, normally have the request-to-send (RTS) circuit to the DSU switched off. The control station begins operation by switching the RTS on and waiting for the clear-to-send (CTS) circuit from the DSU to be switched on. When the CTS is switched on, the control station begins transmitting addressing characters which are to be delivered to all outlying stations via the digital channels and MJUs. These addressing characters indicate to customer-supplied communication controllers at outlying stations whether the data message is destined for them. The station selected then switches on its RTS lead to its DSU, waits for CTS from its DSU, and responds to selection with a message to the control station. Communication then continues between these two as if they were on a two-point private channel. When communication is completed, the control station and outlying station exchange completion notices and switch off their RTS leads, and the circuit returns to its idle state.

The implications of this type of operation are that the idle state of the multipoint DDS circuit has circuit-idle characters generated by DDS office channel units (OCUs) in all channels. When the control station is transmitting or receiving from an outlying station, data is flowing along those channels but idle codes remain in all other branches. The idle codes, however, are suppressed by the MJUs in the circuit as described previously, thus there is no interference from other branches. Moreover if transmission errors occur on idle branches, the errors are suppressed unless they happen to convert the idle code into a data byte. The advantage of this type of operation has a corresponding disadvantage, namely, that this security has been paid for by sacrificing the time required to start up communications (i.e., switching RTS to on, waiting for CTS, etc.).

An alternative method for operating a multipoint DDS circuit is to require all stations (including the control station) to normally have their RTS lead switched on and to be transmitting a constant steady series of 1's as an idle pattern. In this case, there is no need to wait for CTS, and the control location can begin operation by merely transmitting addressing characters. As described above, all stations receive these characters, one is selected, and that one may immediately begin transmitting to the control station. Of course, the data transmitted from this station is combined by the MJU with data from all other branches; however, the other stations are transmitting a steady series of 1's and, when they are combined with data in the AND gate of the MJU, the data remains as it was on input.

This type of operation clearly has less delay and, consequently, higher efficiency than the one described previously since there is no requirement to switch RTS on and wait for CTS before transmitting. However, the inherent disadvantage is that to the MJU, it appears that all stations are transmitting data at all times, even though most stations are transmitting an idle pattern of steady 1's. Therefore, any transmission errors in any branch of the multipoint circuit, which convert a 1 to a 0, will cause errors to occur and propagate through the circuit to the control location.

The DDS multipoint is designed to be compatible with either type of operation described above. The customer chooses by which method he wishes to operate on the basis of the amount of delay he can tolerate and the degree of protection against transmission errors he desires.

3.3 Delay and efficiency of transmission

An important attribute of any multipoint service is the absolute time delay introduced by the communications system. This delay is important because the basic motivation behind multipoint service is one of providing an economical means for connecting several lightly loaded stations to a single circuit. Obviously, wasting more time in establishing a connection on a multipoint circuit leaves less time available for data communication, and, consequently, fewer stations may be connected to a single circuit.

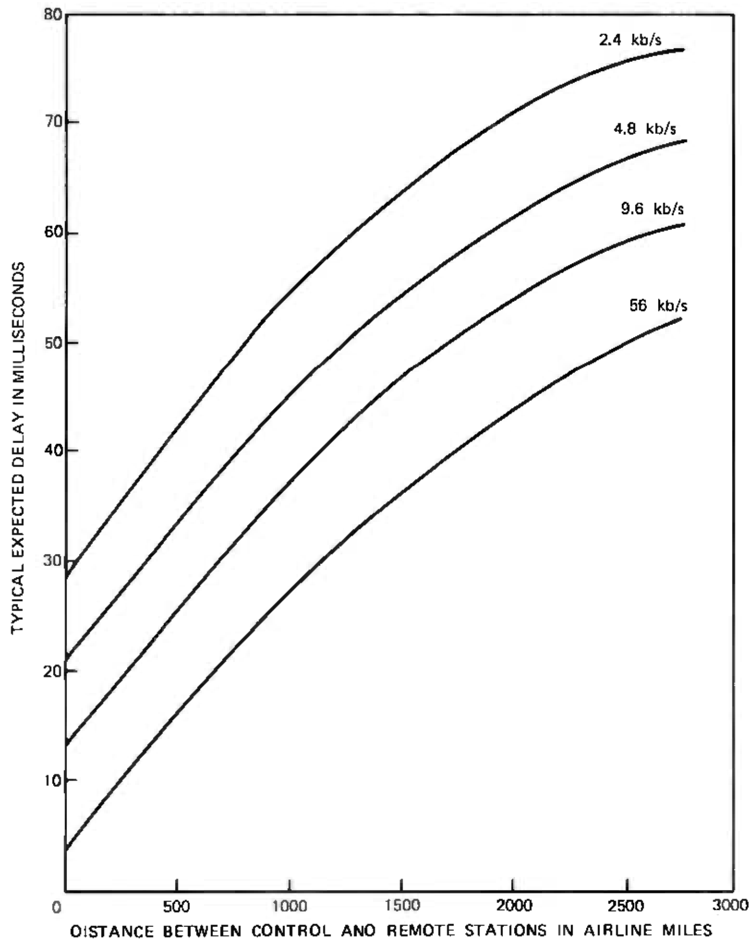

Fig. 1—Typical polling response delays for DDS circuits.

The basic elements of the DDS that introduce delay are the rate at which electrical signals propagate through the transmission media and the built in bit and byte delays in DSUs, OCUs, MJUs, and digital multiplexing equipment. For any specific circuit, these delays can be accurately predicted from knowledge of the distances involved, the equipment through which the circuit is routed, and the actual transmission speed of the circuit.

Figure 1 provides an indication of polling response delay for DDS multipoint circuits as a function of the distance between the control station and the outlying station. This figure is based on the assumption that the customer's master or control station continually keeps its RTS circuit switched on, and transmits polling characters to the remote station. The remote station then switches its RTS on, waits for a CTS indication from the DSU, then transmits a response. Therefore, the delays encountered are one round-trip propagation delay plus one start-up delay at the remote DSU.

When using existing analog data sets, the time delay between RTS and CTS in itself generally exceeds the delays indicated on Fig. 1 by significant amounts. It should be clear, therefore, that similar curves for analog circuits would generally show DDS to have less delay and, consequently, DDS would permit the user to operate his circuit more efficiently than he could using analog. However, the user must be aware that these curves are representative of typical delays and do not portray what may be encountered in any individual case. The specific delays encountered are dependent on the exact routing of a particular channel. Moreover, the user should be aware of the possibility that the delays he experiences on an individual circuit may change due to internal network rearrangements or due to network trouble conditions.

IV. ACKNOWLEDGMENTS

Original contributors to the information in Section II, which describes operation of the Data Service Unit and the Channel Service Unit, are W. D. Farmer, J. R. Klepper, and A. L. Pappas.

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

Digital Data System:

Testing and Maintenance

By S. M. FITCH and D. L. RECHTENBAUGH

(Manuscript received July 12, 1974)

Reliability and maintainability are important aspects of the service objectives for the Digital Data System. Consequently, maintenance planning was an essential element in the DDS development. Maintenance features provided by the system include in-service performance monitoring, protection switching, comprehensive alarms, and the means for rapid fault isolation and repair.

I. INTRODUCTION

The service objectives for the Digital Data System described in the preceding article[1] represent substantial improvements over the performance and availability of existing data services. Since the DDS exists in the present telephone plant environment, it is subjected to the same kinds of random interruptions and failures that occur in that environment. Hence, special arrangements and procedures are provided to meet the more stringent service requirements.

To meet the desired service objectives, DDS equipment at the DS-1 (1.544 Mb/s) and higher levels monitors system performance full-time, with manual or automatic switching to standby equipment in the event of a failure.* Alarms on DDS equipment alert craft personnel to system failures. Many new features in the system permit rapid sectionalization of troubles on a one-man basis and rapid identification and replacement of defective units.

Since the DDS utilizes existing carrier systems for both exchange-area and long-haul transmission, maintenance planning is compatible with present and planned carrier maintenance and restoration procedures. The reliability estimates cited in the preceding paper have, therefore, included allowances for carrier system failures and restoration.

* The subrate data multiplexer (SRDM), which operates at the DS-0 (64-kb/s) level, also provides performance monitoring and protection switching (Ref. 2).

This paper describes the overall maintenance philosophy incorporated into the design of the DDS and the administrative organizations that implement this philosophy. The trouble detection and sectionalization capability designed into the system as well as restoration and repair procedures are also described.

II. OVERALL MAINTENANCE PHILOSOPHY

The maintenance philosophy for the DDS can be divided into two distinct approaches based upon the level of the digital signal, viz., DS-1 and above and DS-0 and below.[3]

2.1 DS-1 (1.544-Mb/s signal) and above

The 1.544-Mb/s signal utilizes existing (and planned) transmission facilities for both exchange area and long-haul transmission. Maintenance and administration of these facilities for the DDS, therefore, cover existing procedures for maintaining the exchange and long-haul transmission networks. Restoration techniques normally employed to protect voice services minimize outage durations. This includes automatic protection switching of long-haul radio channels or coaxial tubes on cable systems.

Since offices in the exchange area are often unmanned and because of the large number of customers affected by a DS-1 channel failure, performance monitoring and automatic protection switching of DDS T1 carrier lines minimize the number of outages and annual outage duration. Terminal equipment carrying DS-1 level signals on both exchange and long-haul facilities provides performance monitoring and either manual or automatic protection switching. In addition to the protection switching feature, the terminal equipment generates office alarms and indications that are useful in the sectionalization and repair of equipment failures.

Briefly, then, the DS-1-level maintenance philosophy is characterized by performance monitoring and protection switching features, use of alarms for trouble sectionalization, and reliance upon existing facility maintenance arrangements.

2.2 DS-0 (64-kb/s signal) and below

At the DS-0 and lower data rates, equipment is modularized on a customer-circuit basis so that a failure usually affects only one customer. Therefore, redundancy and/or protection switching are not provided except in the SRDM, common power supplies, and some common timing distribution circuits where a failure affects many customers. Instead, the ability to alert the customer to service outages and the means for rapid one-man sectionalization of failures are provided.

Testing a customer channel can be accomplished on both an in-service or out-of-service basis, depending upon the type of trouble.

In summary, the maintenance philosophy for DS-0 and lower-level signals is characterized by a high degree of test access and sectionalization capability.

III. ADMINISTRATIVE RESPONSIBILITIES

Since DDS service uses some of the same transmission systems used for voice service and encompasses customer stations in widely separated locations, a number of different organizations have administrative responsibilities for service maintenance. Figure 1 depicts typical equipment configurations* for several offices providing DDS service in a metropolitan area. The figure also indicates the areas of responsibility exercised by the following three administrative organizations charged with maintenance of DDS service:

(*i*) Serving test center (STC)—Responsible for all customer trouble reports and sectionalization of troubles at DS-0 and lower data rates.

(*ii*) T-carrier restoration control center (TRCC)—Responsible for restoration activities on DS-1 level channels in the metropolitan area.

(*iii*) Regional operations control center (ROCC)—Responsible for restoration activities on long-haul facilities.

The responsibilities of each of these organizations as they relate to DDS maintenance are covered in more detail in the following sections.

3.1 Serving test center

The STC is responsible for customer service on an end-to-end, individual-circuit basis. Every DDS circuit appears in at least one STC located at a hub office. From the STC, DS-0 level signals may be transmitted, received, or monitored on any customer circuit in either direction of transmission. The primary maintenance activities of an STC are receipt of customer trouble reports, circuit testing, restoral of service outages, and various administrative functions.

The STC performs three types of circuit tests: performance evaluation, in-service monitoring, and fault localization. Performance evaluation involves measuring the error performance of a circuit and is

* Most equipment shown in Fig. 1 has been described in detail in other papers in this issue (Refs. 2, 4, and 5). Brief descriptions of those units not described elsewhere appear in Appendix A of the paper entitled "System Overview."

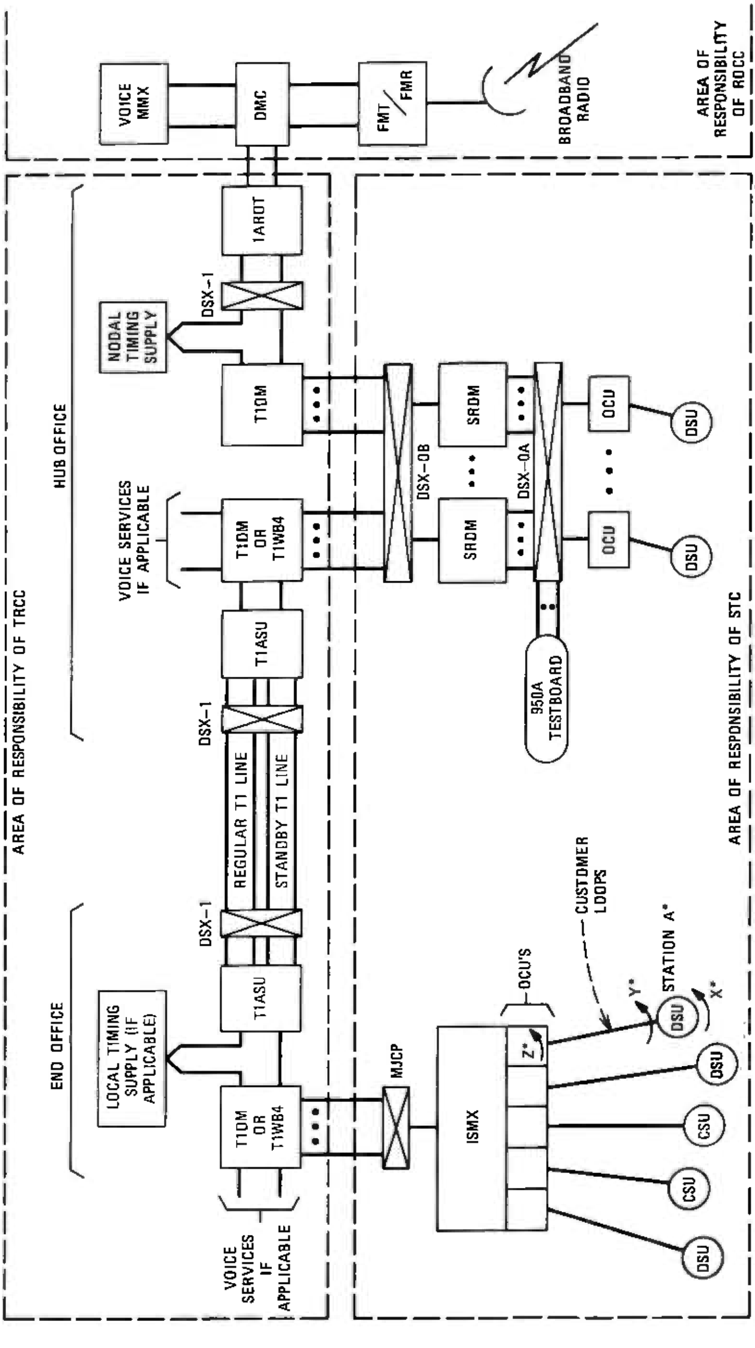

Fig. 1—Administrative responsibilities.

normally done before a circuit is turned over to a customer initially, after a service rearrangement, or after a service outage has been restored. In-service monitoring of an individual circuit is done on a byte-by-byte basis. The existence of certain repeated control bytes indicates specific circuit failures, thus aiding in the recognition and localization of failures. Fault localization permits one man to isolate faults to the network or to the station apparatus or loop facilities in either his local metropolitan area or a distant metropolitan area. These tests can be made on either two-point or multipoint circuits. Fault location within the network usually requires the cooperative effort of other STC or central office craft forces.

Customer service is restored after a fault has been isolated. Restoration activities include referral of trouble indications to other organizations such as the TRCC or ROCC, requests for dispatch of central office or station craft forces, and patching at the DS-0 level. Other functions include the responsibility for coordinating these activities and verifying that a circuit is again operational.

The administrative functions of an STC include administration of the DS-0 cross-connect, maintenance of individual customer records, and summary and reporting of DDS service results.

3.2 T-carrier restoration control center

The growth of T-carrier systems in many metropolitan areas has resulted in the implementation of administrative centers known as TRCCs. These centers have overall responsibility for administering the restoration of T-carrier facilities within a metropolitan area. Where possible, the restoration is on a terminal-to-terminal basis using "backbone" lines (maintenance spare lines) to minimize outage time. Sectionalization and fault-locating activities are the responsibility of the central office craft forces.

Since the T1 lines carrying DDS service are protected from most service interruptions by the T1ASU, the main function of the TRCC with respect to DDS service is to restore the failed line so as to minimize the interval during which the channel is being operated in an unprotected mode.

In the event of a service interruption, the TRCC advises the STC of the status of T1 systems carrying DDS service so that customers may be properly appraised of the situation. The TRCC also ensures that non-service-affecting DDS equipment failures are repaired promptly to minimize the duration of the unprotected mode of operation. This latter function will become increasingly important with the implementation of centralized alarm reporting arrangements and the resultant uncovering of many offices.

Alarms generated by the 1A radio digital terminal[6] can be used by the TRCC in analysis of terminal failure indications. In addition to alerting the central office craft personnel of 1ARDT equipment failures and the need for initiation of a switch to the spare terminal, these alarms can alert the TRCC of long-haul failures, thereby aiding in the interpretation of sympathetic alarms* generated within the metropolitan area. In the absence of a TRCC, these functions are performed by a terminal office.

3.3 Regional operations control center

Present long-haul broadband transmission facilities are administered by regional operations control centers. The primary function of these centers is to select and control the execution of plans to restore broadband facility failures. These centers are well established and interact with the STC and TRCC to provide information on the status of broadband channels carrying DDS service.

IV. CUSTOMER SERVICE

Customer service in the DDS is characterized by the service objectives described in Ref. 1. These service objectives are:

(*i*) Quality—Average of at least 99.5-percent error-free seconds.
(*ii*) Availability—Long-term average of at least 99.96-percent channel availability.

To achieve these objectives, the DDS provides in-service monitoring and either automatic or manual protection switching, equipment alarms, and rapid isolation and restoral of customer-reported troubles.

4.1 Service protection

As mentioned in Sections 2.1 and 2.2, the ability to protect DDS channels from incurring appreciable outage time is provided for all signals at the DS-1 level and above and usually for those DS-0 level signals that affect many customers.

4.1.1 Equipment

The three DDS equipments that provide in-service monitoring and automatic protection switching are the T1 data multiplexer (T1DM), the T1 data-voice multiplexer (T1WB4), and the SRDM. In each case, a performance monitor continuously verifies the operation of the equipment.[2] In the event of a detected fault, a protection spare is

* Sympathetic alarms are simultaneous alarms generated when a failure is detected in other equipment on the same circuit as the failed equipment.

automatically switched in place of the faulty unit and an office alarm is generated. The level of performance that results from fault detection and automatic protection switching is expected to satisfy the overall performance objectives cited above.

The 1ARDT provides a DS-1 data channel on existing radio systems for DDS as well as other services. It uses signal level and format monitoring to detect terminal failures and provides a manually initiated protection switching arrangement for service protection. If the 1ARDT location is not manned, the alarms and protection switching controls can be transmitted to a remote manned location.

4.1.2 Facility protection

Long-haul facilities carrying DDS service utilize the same protection-switching arrangements as are provided for voice services. Specifically, broadband radio systems provide the capability of automatically switching to a protection channel in the event of valid signal loss or fading. At the ends of each message unit radio link, automatic switching arrangements are provided to protect against failure of the wire line entrance link or FM transmitter or receiver. Digital Data System channels utilize these same protection switching arrangements. Similar arrangements are inherent in the design of all higher-level transmission systems.

In the exchange area, automatic protection switching arrangements are not normally provided for voice services. Therefore, a protected T1 line arrangement has been developed for the DDS. This arrangement provides a dedicated T1 line as a standby for each DDS line on a terminal-to-terminal basis. The lines are double-fed at the transmit terminal and switched at the receive terminal. Bipolar violations and pulse absences are the criteria for determining whether the regular and standby lines are working satisfactorily. Office alarms are generated in a failure of either or both T1 lines.

4.2 Trouble detection

Digital Data System central office equipment provides alarm information for detecting and sectionalizing failures. This information is presented locally within the office in the form of major and minor audible alarms as well as visual aisle pilots and status indications on the failed equipment. The major audible alarm is used only for a service outage and indicates to the craft forces that prompt corrective action is required. The minor audible alarm indicates that a failure has occurred that places the service in jeopardy of an outage. The visual aisle pilots locate the equipment generating the alarm, while the status indications denote the nature of the failure.

DS-1 level equipment (and the SRDM at the DS-0 level) also provides a set of status indications for transmission to a central location such as the TRCC. These indications can be used to rapidly sectionalize troubles from the central location, especially when failures occur in unattended offices.

Because of the large number of customer circuits affected by a DS-1 channel failure, extensive trouble detection capabilities are provided in DS-1 level equipment. Central office craft forces are alerted via office alarms to such failures as loss of input signal, clock or power failure, or high error rate. As a result, those troubles reported by the customer to the STC that have not already been brought to the attention of the central office craft forces normally relate to faults of individual DS-0 channels, loop facilities, or station apparatus.

V. TROUBLE ISOLATION

5.1 Alarm analysis

As described above, DS-1-level signal failures are detected and sectionalized by analysis of the office alarms. In nonservice-affecting failures, the alarms are confined to the failed equipment or that unit which detected the failure condition. These alarms are readily identified and do not require extensive analysis. However, in a service-affecting failure, sympathetic alarms may be generated. These alarms must be analyzed to determine the nature and location of the failure. As an example, Fig. 2 depicts a typical exchange area DS-1 channel between two offices. As shown, the channel consists of DDS equipment (T1DMs and T1ASUs) located in the terminal offices (A and Z) and interconnected by T1 lines. The regular and standby T1 lines are part of the protection switching arrangement provided by the T1ASU and usually employ diverse routing to ensure service protection. The restoration backbone line is a manually patched T1 line which appears at the DS-1 cross-connect (DSX-1) in each terminal office. It is administered by the TRCC and is used to temporarily replace a failed T1 line. The backbone line normally restores T-carrier voice services and is not dedicated to the DDS. The T1 lines pass through a number of intermediate offices (B through J) and are available at office repeater bays for testing during trouble sectionalization. However, signal failure alarms are only generated at the terminal offices. Figure 3 is an example of the alarm analysis procedures used by craft forces in a terminal office to localize troubles and restore service on a DS-1 channel such as that shown in Fig. 2. As shown in the figure, certain failures, such as a DS-1 signal failure, result in alarms from both the T1DM and the T1ASU, while other types of failures cause alarms in only one of the

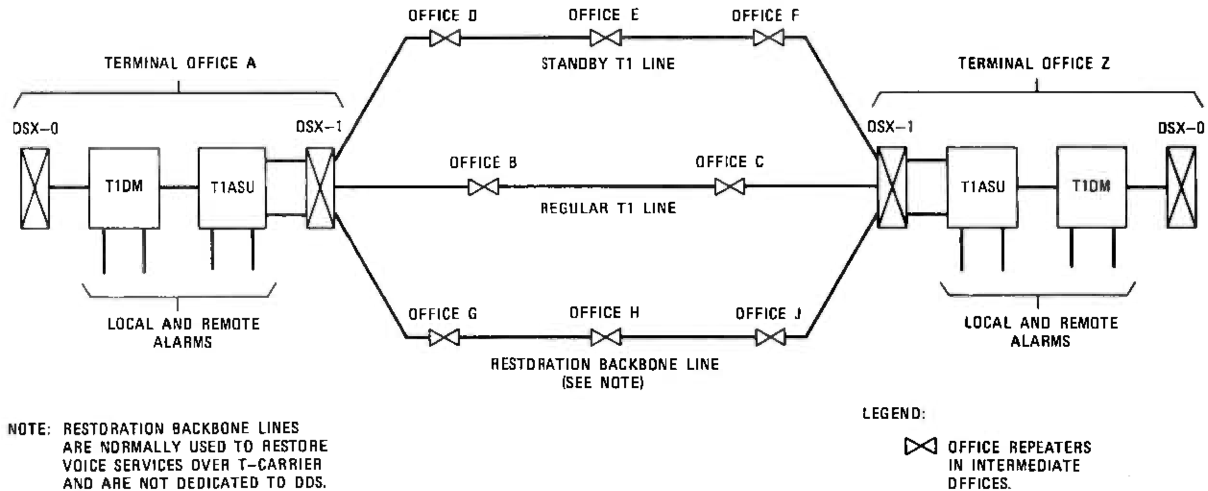

Fig. 2—A typical exchange area DS-1 channel.

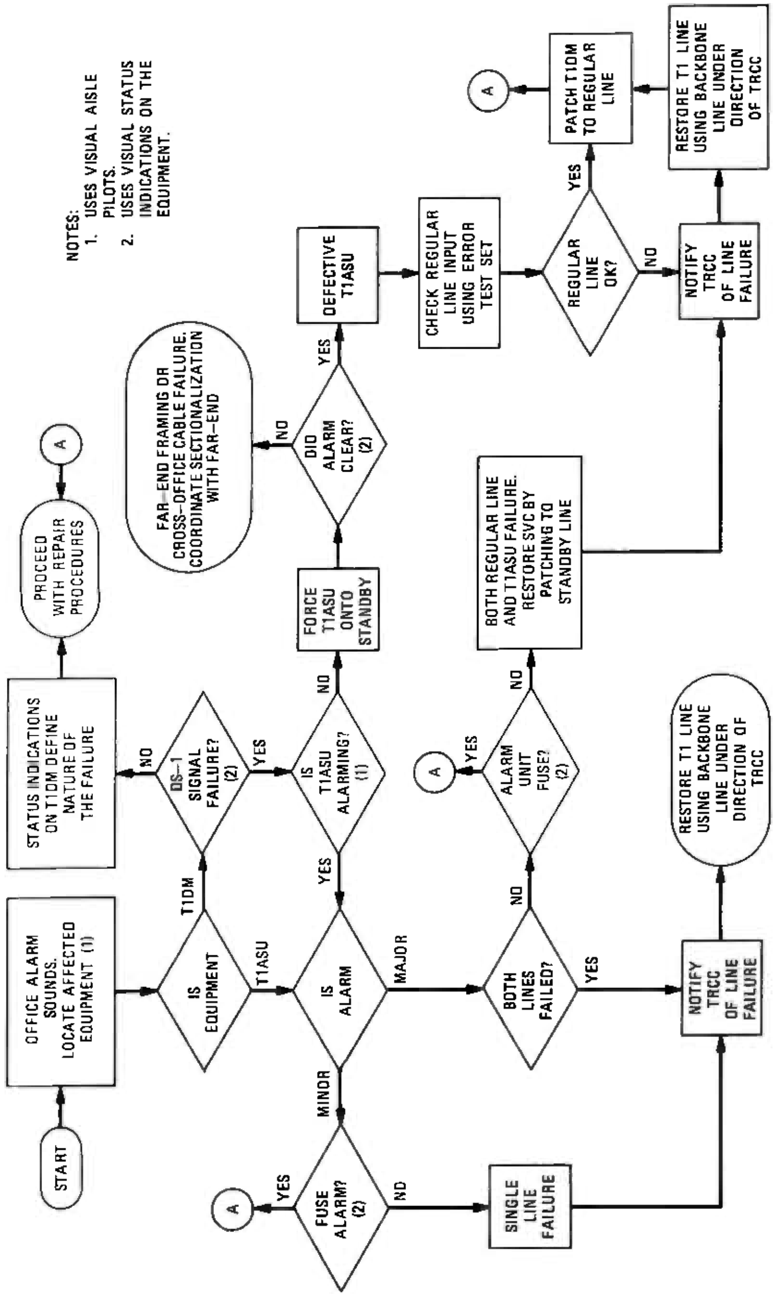

Fig. 3—DDS alarm analysis procedures.

units. In addition, some types of failures require coordination with the distant terminal office or the TRCC. It should be noted that the sectionalization procedures require minimal circuit testing to sectionalize failures to the defective unit and effect restoration. The alarm analysis described above can be conducted either by the local office craft forces using the office alarms or by a central location, such as the TRCC, using status indications at a remote manned location. Similar procedures are required to sectionalize other possible alarm conditions generated by DS-1-level equipment.

5.2 *Customer-reported troubles*

As noted earlier, all customer trouble reports are received by an STC. An STC is equipped with 950A testboards, supplemental jack bays, and a DS-0 cross-connect (DSX-0), each of which is specifically designed for use in the DDS. The 950A testboard and supplemental jack bays provide test access to individual customer circuits, while the DSX-0 permits interconnection of DDS equipment at the 64-kb/s level. Each testboard contains two newly designed data test sets, the digital transmitter and the digital receiver. The transmitter operates at the DS-0 single channel level and the receiver operates at both the DS-0 single and multiplexed channel level. In addition to their application in the STC, these test sets are used as portable units in DDS equipment areas.

When a customer trouble report is received at an STC, the customer circuit is monitored in both directions of transmission on an in-service basis for specific repeated control bytes. The existence of these control bytes indicates that certain fault conditions are present within the network. An example is the absence of a DS-1 signal on the long-haul portion of the circuit. This fault results in the repeated generation of a fixed control byte by the receiving T1DM. Whenever repeated control bytes suggest a network fault, further testing is required to isolate the fault. Between STCs, testing of DS-0 signals requires the cooperative efforts of craft forces from each STC. Within the metropolitan area, DS-0 level testing usually involves signal tracing and the cooperative efforts of STC and central office craft forces. The digital transmitter and receiver are used as portable units at the DDS equipment frames for this purpose. In local offices, jack and connector panels provide jack access in testing (as well as the interconnection of DDS equipment).[7] In addition, all input-output signals at the 64-kb/s level on all DDS equipment are available on individual circuit-pack face plates.

If no indication of a network failure is provided by monitoring, then loopback tests are made of the data service unit (DSU) or channel service unit (CSU) and the office channel unit (OCU).[4] The STC can remotely test these units by transmitting alternate loopback control bytes and pseudo-random data. During loopback, data received from the STC are retransmitted back to the STC and error performance measurements are made. The tests can be made for every station on a customer circuit, whether located in a local or distant city or on a two-point or multipoint circuit. On a multipoint circuit, a station is selected for testing by means of the multipoint signaling unit (MSU) described later.

As an example of loopback testing and its use in fault isolation, consider station A in Fig. 1. Three loopbacks are associated with this station, one at the DSU-customer interface, one at the loop-DSU interface, and one at the OCU-loop interface. These are labeled X, Y, and Z, respectively, on Fig. 1 and are called DSU loopback, channel loopback, and OCU loopback. First, DSU loopback is attempted, followed by error performance measurements. If these tests are successful, then the customer circuit between the customer interface and the STC is satisfactory. If these tests are unsuccessful, then channel loopback is attempted. If these tests are successful and a preceding DSU loopback is unsuccessful, then a fault exists in the DSU. If channel loopback is unsuccessful, then OCU loopback is attempted next. If these tests are successful and a preceding channel loopback is unsuccessful, then a fault exists on the loop. If these tests are unsuccessful, then a fault exists between the OCU-loop interface and the STC. By using these techniques, STC personnel can rapidly isolate a circuit trouble to the DSU (either local or distant), to the loop (either local or distant), or to the network. Similar procedures are used for CSUs.

Loopback testing for multipoint circuits[1] requires special procedures at the STC. Consider the typical multipoint circuit in Fig. 4. This circuit consists of seven stations and three multipoint junction units (MJUs). The MJUs are located only in hub offices where all input/output ports of the MJU appear at a 950A testboard. It is apparent that, if a loopback command were transmitted downstream from the St. Louis STC on the main channel (the unnumbered input/output port of the MJU), all stations B through G would loop back simultaneously. The result would be garbled data at the digital receiver of the St. Louis STC. Consequently, before loopback testing on a multipoint circuit can begin, a specific station must be selected for testing. Using the multipoint signaling unit, an STC selects one branch (a numbered input/output port) of each downstream MJU for data transmission and blocks all other branches. The selection process involves a prescribed control dialogue in the DS-0 channel between the MSU and the MJUs. For example, an MSU in St. Louis can select branches one, four, and two in St. Louis, Chicago, and Boston, respectively, to set up an equivalent two-point circuit from the MSU to station G. Stations B through F

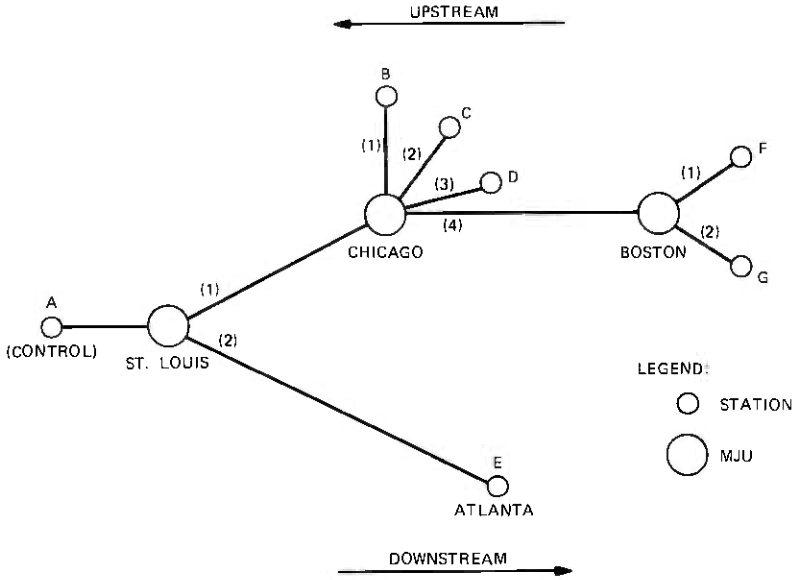

Fig. 4—A typical multipoint circuit.

are blocked and cannot transmit data to or receive data from the MSU. Now loopback tests for station G can be made as described above. Each station in turn can be tested in a similar manner.

If the control dialogue between an MSU and any MJU of a multipoint circuit fails, the affected STCs must determine whether the DS-0 channel between STCs or the MJU itself has failed.

VI. SERVICE RESTORATION

A service outage in DDS is restored by either of two techniques, depending upon the type of fault. The first is to manually patch around the fault and the second is to replace defective units.

6.1 Below DS-0

Station apparatus (DSUs or CSUs), loop facilities, and OCUs all operate below the DS-0 level. As described above, a fault can be isolated to each of these units from the STC. Once isolated, the method of restoration is different for each.

Station apparatus is restored by dispatching craft forces to the customer's premises. Service is usually restored by replacing the DSU or the CSU.

A loop facility is usually restored by isolating the fault to a cable segment with convenient access at each end. Restoration is then accomplished by changing to a spare cable pair if available. If a spare is not available, the fault must be repaired. Since the DDS uses the standard loop plant, existing fault location and repair procedures are normally followed.

An OCU is restored by replacing defective circuit packs, power supplies, or fuses. The defective unit is isolated through analysis of bay alarms and use of portable test sets.

6.2 DS-0

Integral subrate data multiplexers (ISMXs),[4] SRDMs and associated performance monitors, and MJUs all operate at the DS-0 level. Restoration of this equipment again involves replacement of circuit packs, power supplies, or fuses. Defective units are isolated by the craft forces through the analysis of bay alarms, the use of the portable test sets, and, in the case of the SRDM, status indications. The status indications are a combination of indicator light-emitting diodes (LEDs) and LED digit readouts that specifically pinpoint defective units.[2]

6.3 DS-1 and above

As indicated in Section 2.1, service outages on DS-1 and higher level channels are normally encompassed by the procedures used to restore existing carrier systems. In the exchange area, the TRCC is notified of any T1 line outages by the terminal office and coordinates the facility restoration activities. Restoration of broadband facilities is coördinated by the ROCC.

As high-capacity routes between offices and cities in the DDS network are established, they will utilize higher-level transmission systems and will make restoration of individual DS-1 channels feasible by alternate routing by patching at the DSX-1. This alternative will, however, only be used when restoration cannot be effected using the normal transmission facility restoration plan.

DS-1 level terminal equipment (T1DM, T1WB4, T1ASU, 1ARDT, local timing supply, and nodal timing supply) is equipped with alarms, status indications, and local test features that enable craft personnel to isolate failures to a particular unit. Restoration is effected by replacing the defective circuit pack, power supply, or fuse. Occasionally, a defective T1ASU must be bypassed by manual patching to restore service while the unit is repaired.

VII. REPAIR

The repair philosophy for all DDS equipment in the central office is replacement of defective units such as circuit packs, power supplies, or fuses. The defective units themselves are usually repaired at a Western Electric Company service center. Other failures such as faults in intraoffice cabling or connectors must, of course, be repaired on site. Station apparatus (DSUs and CSUs) are normally returned to a service center for repair.

Equipment and facilities not designed specifically for the DDS and shared with other services are repaired using existing repair procedures. Examples of these are customer loops, T1 lines, and radio systems. Spare station apparatus, circuit packs, and power supplies are required in sufficient quantities to satisfy DDS service objectives. In the case of unprotected units, service restoration depends upon rapid repair so that spare units are normally available at all central office locations. In the case of protected equipment, spare units are available in quantities sufficient to minimize the unprotected mode of operation and may be kept at a centralized location.

VIII. SUMMARY

Service objectives for the DDS place special emphasis upon the means for adequately maintaining the system. Throughout the development of the DDS, special consideration was given to the incorporation of maintenance features that aid in achieving the overall service goals.

To reduce the probability of a service outage, full-time performance monitoring and either manual or automatic protection switching arrangements are provided on the SRDM and on all equipment operating at the DS-1 level and above. Alarms are provided to alert central office craft personnel so that repair can be rapidly effected. In a service outage, these alarms, together with appropriate sectionalization procedures, permit rapid restoration of the service. The restoration procedures for facility failures are normally those used for voiceband services.

Below the DS-1 level, service outages are usually detected and reported by the affected customer. Provisions have been made in the DDS equipment design to enable the serving test center to rapidly sectionalize equipment and loop failures on a one-man basis. Restoration of service is then effected by isolation and replacement of the defective units.

IX. ACKNOWLEDGMENTS

The authors are pleased to acknowledge the contributions of L. F. Bugbee, S. Caputo, J. N. Daigle, A. V. Gallina, J. J. Mahoney, Jr., G. N. Packard, J. C. Panek, D. C. Rife, D. v. Z. Wadsworth, and

many others, in developing the maintenance plans, arrangements, and equipment for the DDS.

REFERENCES

1. J. J. Mahoney, Jr., J. J. Mansell, and R. C. Matlack, "Digital Data System: User's View of the Network," B.S.T.J., this issue, pp. 833–844.
2. P. Benowitz, S. J. Butterfield, M. P. Cichetti, Jr., and T. G. Cross, "Digital Data System: Digital Multiplexers," B.S.T.J., this issue, pp. 893–918.
3. N. E. Snow and N. Knapp, Jr., "Digital Data System: System Overview," B.S.T.J., this issue, pp. 811–832.
4. E. C. Bender, J. G. Kneuer, and W. J. Lawless, "Digital Data System: Local Distribution System," B.S.T.J., this issue, pp. 919–942.
5. B. R. Saltzberg and H. M. Zydney, "Digital Data System: Network Synchronization," B.S.T.J., this issue, pp. 879–892.
6. K. L. Seastrand and L. L. Sheets, "Digital Transmission Over Analog Microwave Radio Systems," Proc. IEEE 1972 Int. Conf. Commun., *8*, IEEE Cat. No. 72-CHO-622-1-COM, Philadelphia, Pennsylvania, June 19–21, 1972, pp. 29-1—29-5.
7. H. C. Illium, W. B. Lueft, and D. W. Rice, "Digital Data System: Physical Design," B.S.T.J., this issue, pp. 943–964.

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

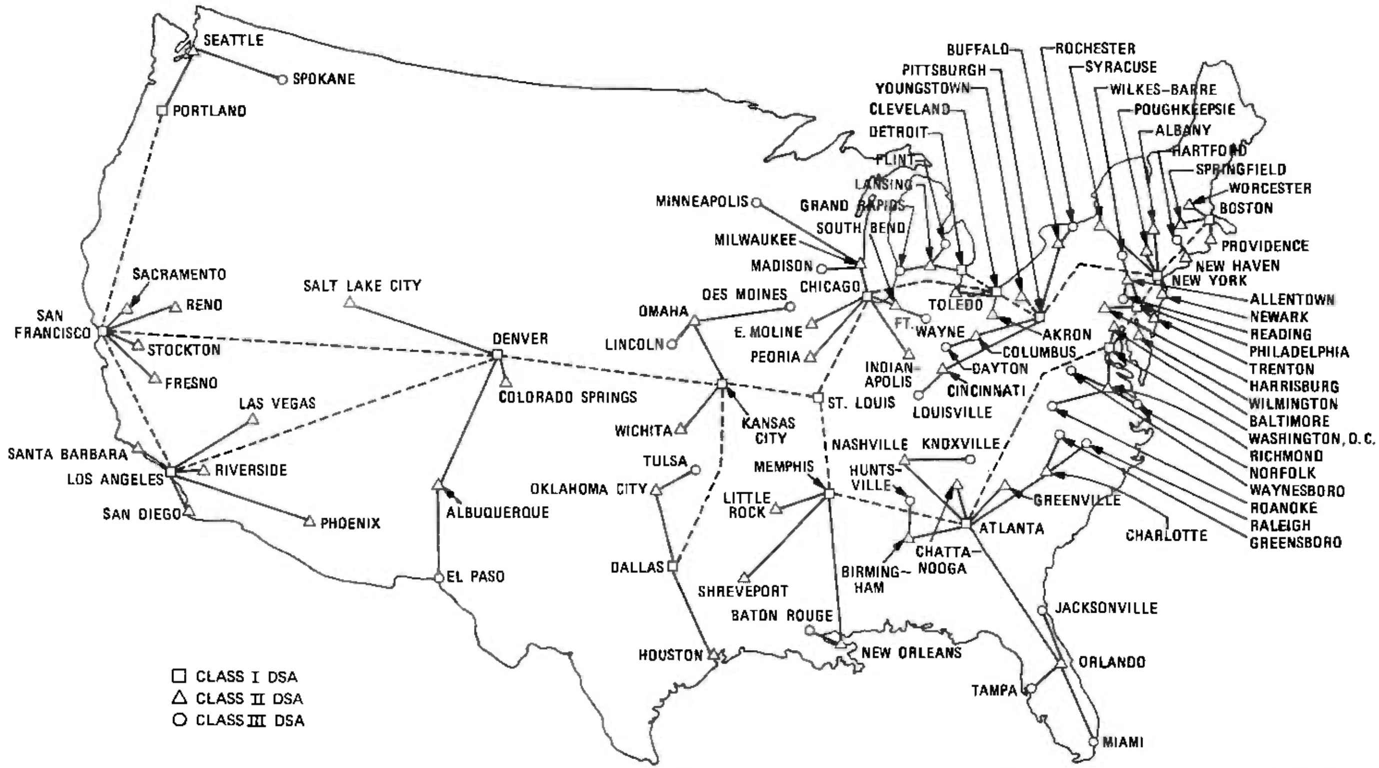

Fig. 1—Projected 96-city network.

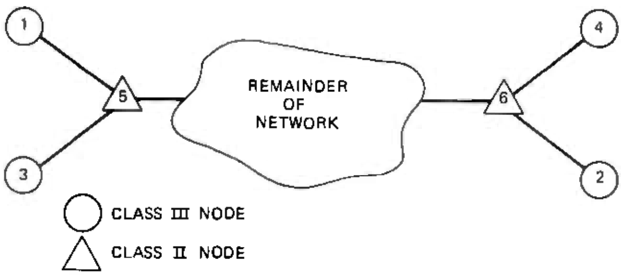

Fig. 2—Multiplexing algorithm example.

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.

(a)

	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						

(b)

	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						

(c)

	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						

(d)

	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						

(e)

	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						

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

$$24 - n20 \geqq 0$$
$$n = 1.$$

Now determine

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

If $R = 0$ or $R \geqq F_r C_r$, 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 $R \geqq F_r C_r$, 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 $R < F_r C_r$, 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 $R < F_r C_r$ 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

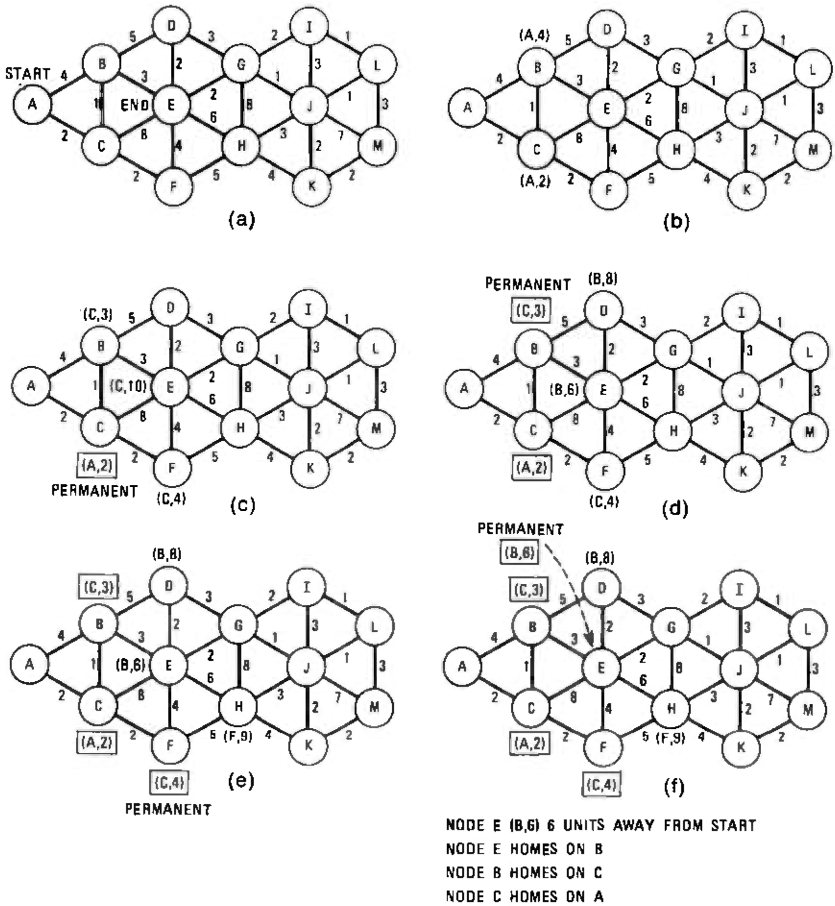

Fig. 4—Dijkstra minimum-distance algorithm.

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

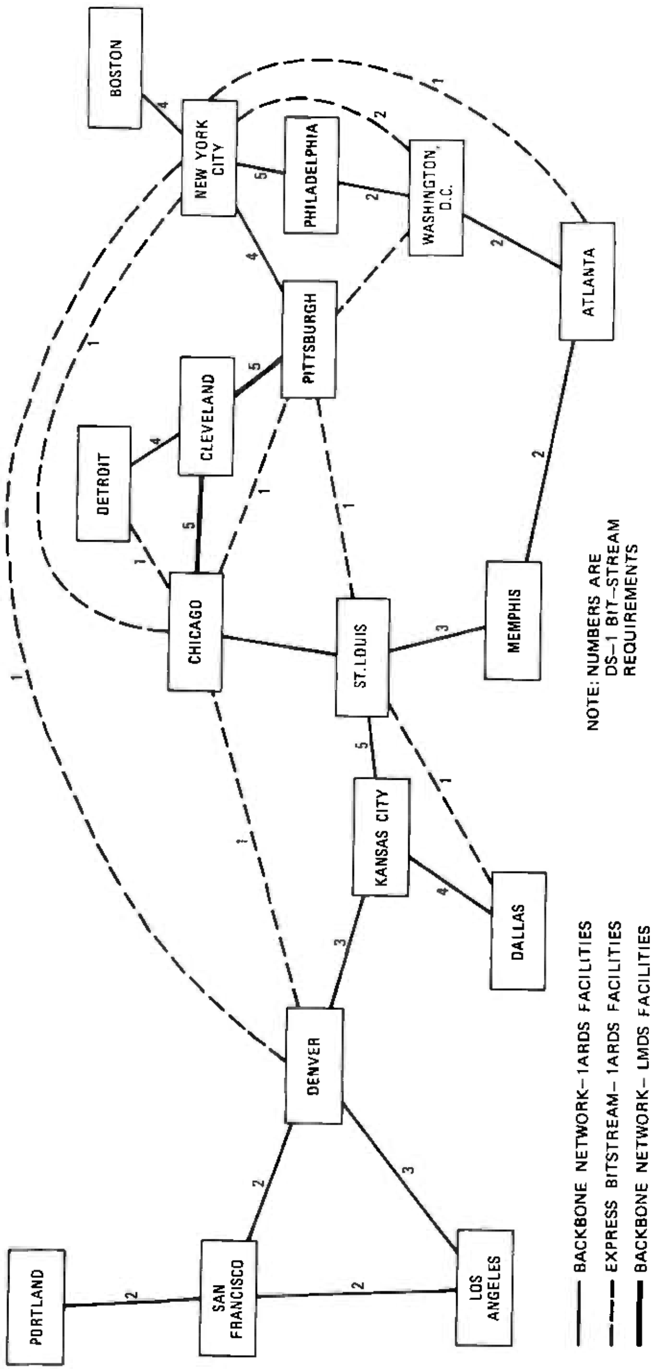

Fig. 5—Example network subset—backbone and express bit streams.

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.

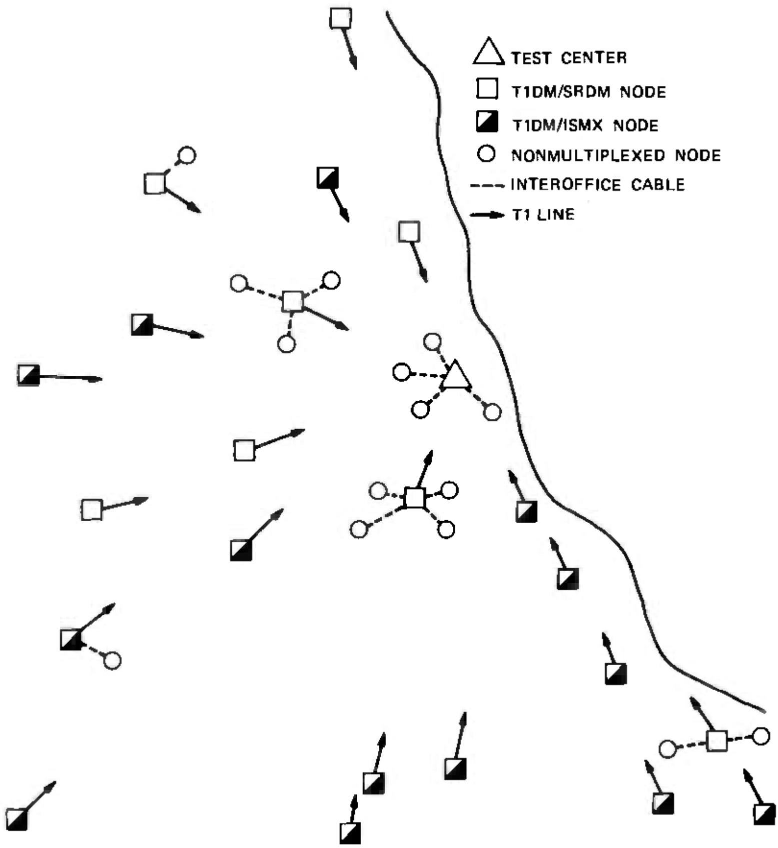

Fig. 6—Typical DSA network.

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

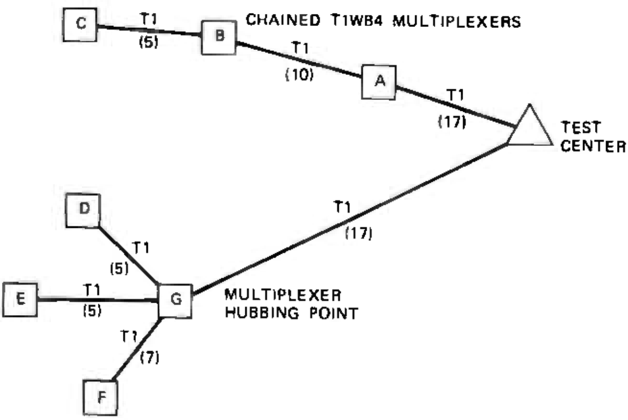

Fig. 7—Digital line sharing.

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-

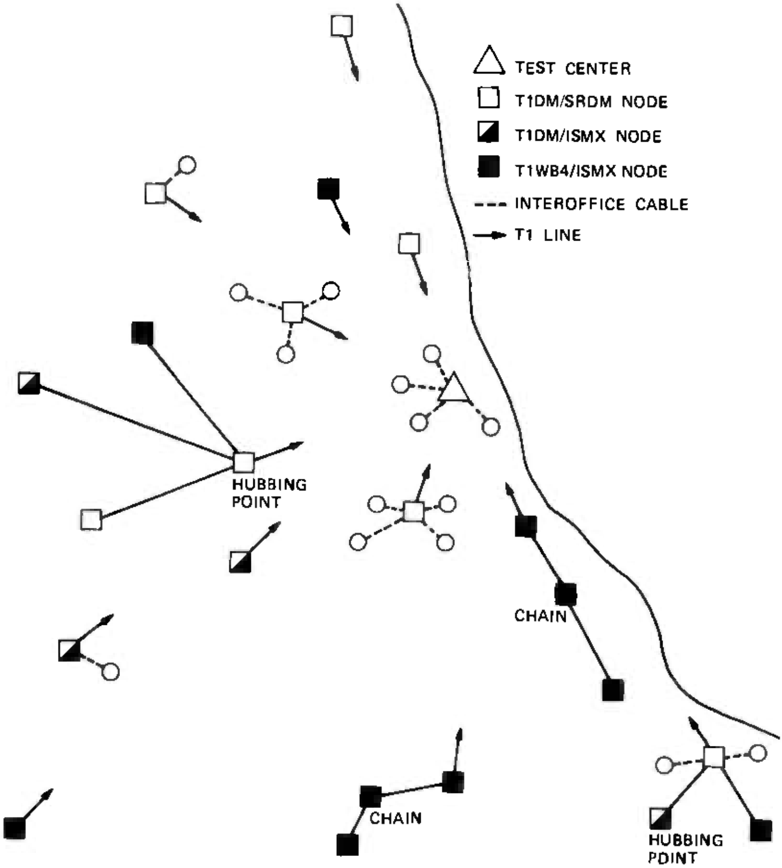

Fig. 8—Typical DSA network for metropolitan area.

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.

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

Digital Data System:

Network Synchronization

By B. R. SALTZBERG and H. M. ZYDNEY

(Manuscript received July 12, 1974)

The Digital Data System is synchronized by means of a network in the form of a master-slave tree. Functioning data links are used for distribution of synchronization signals and timing is recovered from the data stream. The master timing supply sets the frequency for the entire network. Each major node contains a nodal timing supply with an extremely long time constant, which phase-locks to the incoming data signal. The nodal timing supply contains frequency memory and sufficient accuracy to free run satisfactorily for several days if its inputs fail. A local or secondary timing supply is designed for operation in distant nodes of the tree. Each timing supply provides common clock signals to all DDS equipment in the office in which it is installed. All timing supply designs include a high degree of redundancy for reliability purposes.

I. INTRODUCTION

The network for the DDS consists of an interconnected set of digital transmission facilities. At any node of the network, if the average rate of transmission bits leaving a node is not exactly equal to the average rate entering the node, errors occur. Such errors are defined as slips, which means that bits are arbitrarily deleted if the input rate exceeds the output rate, or that bits are repeated or inserted arbitrarily if the input rate is slower than the output rate. The best performance is realized if no slips occur. This requires that every node in the system be synchronized to the identical average frequency and be capable of absorbing delay fluctuations.

The transmission media comprise radio, coaxial cable, and twisted pairs. Each of these admits to small but troublesome variation in delay because of thermal and other effects. The rates of change of such delay variations are equivalent to small frequency disturbances. To be free from slip under these conditions requires two elements: (*i*) reproduction of identical average frequencies at every node, and (*ii*) storage of sufficient data at each point of entry to accommodate the inevitable delay variations of the transmission media. The details of elastic storage to accomplish the second function are discussed in the description of digital multiplexers.[1] The frequency distribution is the subject of this article.

II. SYNCHRONIZATION NETWORK

Distribution of accurate frequency for transmission systems has been the subject of many studies.[2] For DDS, a loosely coupled master-slave system was selected as the appropriate process to meet performance, cost, and administrative requirements. In normal operation, frequency information is transmitted down a topological tree which is a selected subset of the actual data facilities. Such a tree is shown in Fig. 1.

A reference frequency originates the timing signal at a location convenient for other Bell System purposes. This is directly transmitted to a master timing supply that creates a timing format compatible with the remainder of DDS. This signal is transmitted, via the 1.544-Mb/s DS-1 data stream, to major nodes containing nodal timing supplies. Each of these is phase-locked to the incoming data signal with a very long time constant so that the frequency generated at each node is extremely stable and always quite close to the average incoming frequency. The unit of phase controlled at each node is the 125-μs interval that represents a multiplexing frame. Elastic storage in each data-stream input withholds a sufficient number of bits so that data delivered to the node can be held in exact frame alignment

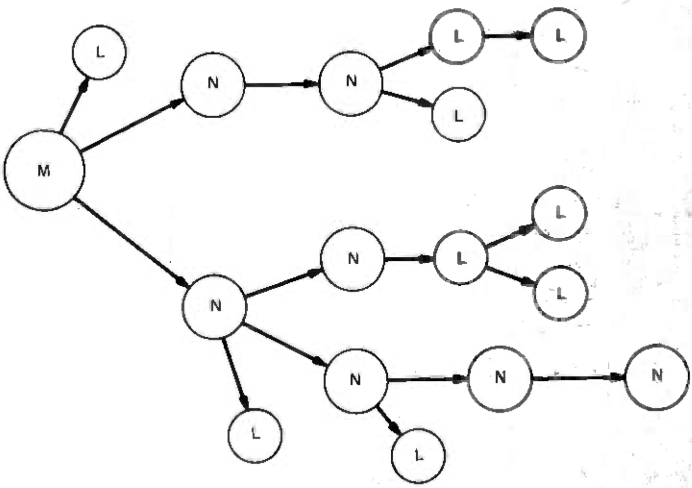

Fig. 1—Synchronization tree of master, nodal, and local timing supplies.

from each source, using the synchronization data stream as reference. This process is repeated at each major node in the tree. Minor nodes, with more limited connectivity, are equipped with phase-locked local timing supplies that are simpler in design and have somewhat more relaxed specifications, but which otherwise perform a similar function. When interconnected, such a system performs slip-free and is an efficient approach to timing distribution across the continental United States.

A primary concern in implementing a system of this sort is the possibility of a temporary interruption of transmission supplying a branch of the synchronization tree. Nodal timing supplies are designed to control the frequency of the internal oscillator by means of a digital proportional-plus-integral phase-locked loop in which the integral is held in a digital memory. This is described in detail in Section 3.1.2. Interface circuitry constantly monitors for the presence of a validly formatted timing signal in the data stream. If it is interrupted, further phase information is inhibited and the oscillator can then run free at the last remembered frequency stored in the integral memory. The interface circuitry automatically returns to normal operation when the outage is ended. If maintenance craftsmen determine that the outage will persist, a spare "hot" interface circuit and alternative transmission path is substituted. If a still more severe outage should occur, the nodal timing supply will run free of any timing input, beginning at its last frequency setting. During this condition, the timing supply relies on the precision of the oscillator to maintain synchronism. The selected oscillators hold frequency to an accuracy of one part in 10^{10} drift per day, which should result in a slip rate of less than one frame each day until the timing source is restored.

The local and secondary timing supplies are designed to operate from two independent, protected, 1.544-Mb/s channels, if they are available, and automatically transfer between signal sources. They are at a low level in the tree and a catastrophic failure causing loss of both input signals is unlikely to cause serious problems. They operate without digital storage of phase and contain only moderately accurate oscillators within their phase-locked loops. The secondary timing supply is physically arranged to permit easy conversion to a nodal timing supply when system growth so requires.

III. TIMING SUPPLY DESCRIPTION

3.1 Nodal timing supply

The nodal timing supply delivers a common set of clock signals to all DDS equipment in the hub office in which it is installed. It maintains synchronization with other hubs by phase locking to the framing bit of an incoming DS-1 signal. When the incoming signal is interrupted or defective, the nodal timing supply free runs at the previous input frequency, appropriately averaged. The oscillator in the nodal timing supply phase-locked loop has sufficient long-term accuracy to permit satisfactory operation even after free running for several days.

A high degree of redundancy is provided for reliability purposes. As shown in Fig. 2, the interface unit, the phase-locked loop, and the output circuit are all duplicated. Units of each pair receive their power from different power supplies. Monitoring and control circuitry is provided to reconfigure the system in response to various outages or manual interventions. A suitable output will be present in the event of trouble.

3.1.1 Interface unit

The interface unit extracts the 8-kHz framing signal from an incoming DS-1 line signal. Each interface unit bridges on a T1 data multiplexer receive line. The two interface units are connected to the same line, to separate lines from the same origin, or to separate lines from different origins, depending upon the desired configuration of the synchronization tree. Only one interface unit is actually supplying a framing bit to the phase-locked loops at any given time. Manual switching between interface units is accomplished by means of a switch on a display-and-control panel. To avoid large phase hits when such a switch is made, a manual build-out switch is present on each interface unit for delaying the output in steps of one-twelfth of a cycle. The build out is adjusted during initial installation to bring the two interface unit outputs into as close alignment as possible.

In addition to the extracted framing bit, each interface unit also provides an output signal that indicates whether or not a valid framing bit is being extracted from the incoming line. A minor alarm is actuated if either interface unit is unable to extract a valid framing bit. The framing bit and the status indication from the selected interface unit are coupled to the phase-locked loops and to the control circuitry.

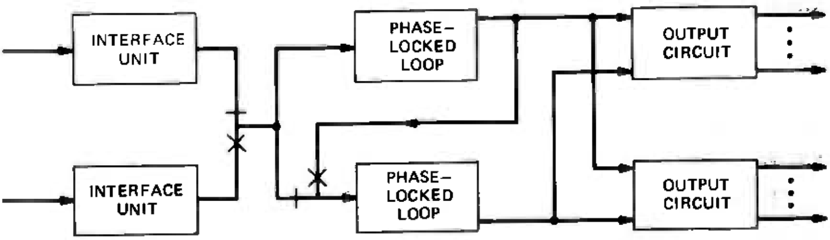

Fig. 2—Nodal timing supply block diagram. Only signal paths are shown.

3.1.2 Phase-locked loop

The phase-locked loop is a digital control system of the proportional-plus-integral type. The proportional-plus-integral phase-locked loop has the interesting property of producing no steady-state phase error in the presence of a steady frequency offset. More importantly, the phase-locked loop will remember its past operating frequency after the input to the loop is removed.

A block diagram of the phase-locked loop is shown in Fig. 3. The principal component is a 39A oscillator, which is a quartz oscillator mounted in a double oven. Its drift rate is less than one part in 10^{10} per day after one day of warm-up. A digital-to-analog converter is built into the oscillator to permit approximately linear fine control of the output frequency in response to a 14-bit parallel digital control input, weighted in a normal binary sequence. The 5.12-MHz output frequency of the oscillator may be varied over a total range of 0.8 part per million by means of this digital control. The least significant bit, therefore, changes the output frequency by a factor of approximately 5×10^{-11}.

The output of the oscillator is divided by a countdown chain to provide a large number of clock and gating signals with periods as long as 8.192 seconds. An 8-kHz and a 512-kHz output are fed to the output circuits. Other outputs from the countdown chain control the phase-locked loop serial arithmetic and provide clock signals to other parts of the timing supply.

In normal operation, both phase-locked loops lock to the output of the selected interface unit. The phase comparator compares the arrival time of this signal with a 4-kHz output from the countdown chain. The phase of alternate input pulses are measured to within 1/320 of an 8-kHz period. The measurement is formed by counting the number of pulses of a 2.56-MHz clock that occur between the start of the input pulse and a transition of the 4-kHz output. The phase comparator is, therefore, of the quantized sawtooth type.[3] The resultant phase measurement is read out serially during the other half of the

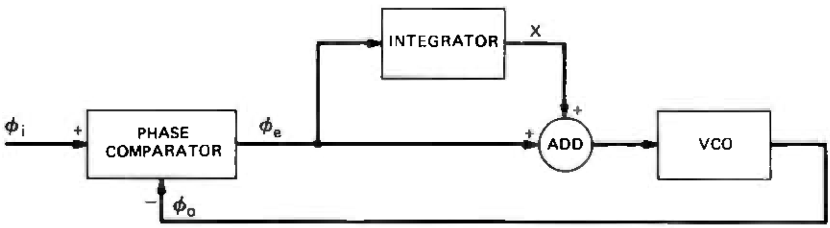

Fig. 3—Nodal timing supply phase-locked loop.

4-kHz signal, during which time no phase comparison is made. During each 8.192-second interval, the average of the 2^{15} phase measurements generated during the interval is calculated. Arithmetic is then performed on this average phase error during each interval to supply the control input of the oscillator. All arithmetic is in serial 2's complement binary form.

The arithmetic is performed in three steps during each 8.192-second interval. The average phase error is first transferred to temporary storage. Next, the average phase error is multiplied by 2^{-15} and added to the contents of an integral register. The integral register maintains a perpetual running sum of the phase error and may be manually reset to zero by means of a pushbutton on the faceplate of a circuit pack. The last step of the arithmetic consists of adding the average phase error, which was temporarily stored, to the contents of the integral register and storing the result back into a register. At the end of the operation, the 14 resultant bits are used to update the frequency control input of the oscillator. The oscillator input is updated only once every 8.192 seconds to reduce the effect on the lifetime of the relay switches in the built-in digital-to-analog converter.

Although the control loop is partially a sampled data system, the sampling interval is much shorter than any of the time constants involved. The summation in the integral register can be approximated very closely by an integral and the resulting open-loop gain function of the phase-locked loop may be approximated by

$$\frac{\Phi_o(s)}{\Phi_e(s)} = \frac{\alpha}{s}\left(1 + \frac{a}{s}\right).$$

The closed-loop gain function is, therefore,

$$\frac{\Phi_o(s)}{\Phi_i(s)} = \frac{\alpha(s + a)}{s^2 + \alpha s + \alpha a}.$$

The quantity α is determined strictly by the gain in the proportional path. In the normal mode described above, a phase error of 1/320 of a cycle at 8 kHz leads to a change of one least-significant bit at the oscillator control input. This, in turn, will change the oscillator output referenced to 8 kHz by a factor of 5×10^{-11}, or 0.4 μHz. α is, therefore, equal to 0.4 μHz divided by 1/320, or 1.28×10^{-4} second^{-1}. The quantity a is determined by the scaling factor of the input to the integral register multiplied by the frequency at which this updating occurs. Since, in the normal mode, the input to the register is multiplied by 2^{-15} and the updating occurs once every 8.192 seconds, a is equal to 3.73×10^{-6} second^{-1}.

The frequency response will show a very slight peaking effect so that some very low jitter frequencies may exhibit very slight amplification. Otherwise, since a is much less than α, the closed-loop response may be approximated by $\alpha/(s+\alpha)$. This is a simple single-pole response with corner frequency at 20.4 μHz for a one-sided noise bandwidth of 32.5 μHz.

The response to an input step in frequency Δf is given approximately by

$$\phi_e(t) \approx \frac{\Delta f}{\alpha}\,(e^{-at} - e^{-\alpha t}),$$

where the phase error is measured as a fraction of a cycle. The phase error, therefore, rises to approximately $\Delta f/\alpha$ with a time constant of $1/\alpha$, which is equal to 7810 seconds or 2.17 hours. The phase error then decays slowly to zero, with a time constant of $1/a$, which is equal to 269,000 seconds or 3.12 days.

After an interruption in the input signal, the phase-locked loop is forced to free run by setting the output of the phase comparator to zero. The output frequency of the oscillator in the free-running mode is then simply equal to α times the output of the integral register, which remains unchanged. While the input is present, the output of the integrator is given by

$$x(s) = \frac{a}{s}\,\Phi_e(s) = \frac{as}{s^2 + \alpha s + \alpha a}\,\Phi_i(s).$$

The corresponding frequency may therefore be approximated by

$$f_x(s) = \alpha x(s) \approx \frac{a}{s+a}\,\frac{\alpha}{s+\alpha}\,s\Phi_i(s).$$

This is equivalent to the tandem combination of two simple single-pole low-pass filters whose corner frequencies are at 0.593 μHz and 20.4 μHz. The first of these filters clearly dominates. One particular jitter frequency, which must be accounted for, is that due to daily variations in path length. Since this is a jitter frequency of 11.6 μHz, the daily jitter is reduced by a factor greater than 20 as far as the output of the integrator is concerned. When the input to the loop vanishes, the loop free runs at a frequency equal to the input frequency prior to the time of free running convolved by an exponential of time constant equal to 3.12 days.

Due to the extremely narrow filtering action described above and to previous line jitter, it is expected that the frequency offset immediately after free running should be less than five parts in 10^{11}, assuming that the frequency of the previous node is perfect. In addition, there is an initial frequency offset of up to five parts in 10^{11} owing to quantization effects. The initial frequency offset of a nodal timing supply immediately after it begins to free run, therefore, is less than one part in 10^{10}. In addition, the oscillator begins drifting at a rate of up to one part in 10^{10} per day. The nodal timing supply may, therefore, free run for up to 13.5 days before the frequency error is equal to the rate of one 8-kHz frame slipped per day. The accumulated phase drift will not reach $\frac{1}{2}$ cycle of an 8-kHz interval until 2.93 days have elapsed. Therefore, assuming zero initial phase error, the phase-locked loop will exhibit no slips if any outage in the input signal is restored before 2.93 days. However, whether or not any customer data will be slipped depends upon the initial fill of the T1 data multiplexer elastic store.[1]

The above parameters are unsuitable for initial lock-in of the loop during installation, or for restoring the loop after certain outages, owing to the narrow lock range and the long time constants involved. For this reason, a fast-start mode has been introduced which is manually entered by operating a switch on a faceplate of a circuit pack. Operation of the switch modifies the countdown chain in the phase-locked loop so as to change the gating waveforms used in the serial binary arithmetic. In the fast-start mode, the average phase error is multiplied by 2^6, thus increasing α to 4.1×10^{-3} second^{-1}. The input to the integral register is further increased by a factor of 2^9, thus changing a to 1.9×10^{-3} second^{-1}. The loop response is changed to a slightly under-damped one with a damping time constant of 488 seconds. Therefore, the loop may be brought into close lock in less than an hour in the presence of large phase or frequency errors.

3.1.3 *Output circuit*

Each of the output circuits accepts an 8-kHz and a 512-kHz waveform from each of the phase-locked loops and generates the composite waveform described in Section IV, Clock Distribution. The two output circuits are coupled so that at any given time both of them receive their inputs from the same phase-locked loop. When a defect is detected in the 8-kHz phase-locked loop output, the output circuits will switch their inputs to the other phase-locked loop and remain there until trouble is detected in that loop.

Sufficient fanout capability is provided and may be expanded to drive an arbitrarily large number of using bays. Each using bay receives its signal over two lines, one originating from each of the output circuits.

3.1.4 Monitoring, control, and alarms

Self-checking circuitry has been built into all major components of the timing supply. Any detected failure will generate at least a minor alarm. If trouble is detected in both phase-locked loops, in both output circuits, or in a combination of two or more power supplies that would render the system inoperative, than a major alarm is generated. An indicator light is provided to caution central office personnel that some manual control is in an abnormal condition, such as the fast start mode, a loop inhibited, or the manual free-run mode.

A built-in digital phase meter is included for manually checking the phase difference between any pair of 8-kHz waveforms in the timing supply.

Built into the nodal timing supply is a control algorithm that controls the input and output of the phase-locked loops in response to certain monitoring circuits. Associated with each phase-locked loop is a slip detector, which looks for an out-of-lock condition in the phase-locked loop. This is done by monitoring the phase difference between the phase-locked loop input and its 8-kHz output. If the phase difference proceeds from a region of moderate positive error to a region of moderate negative error, or vice versa, then the slip detector will generate a slip indication. This indication must be reset manually. In addition to the two slip detectors, a tracking detector is included which monitors the phase difference between the two phase-locked loop outputs and generates an output if the magnitude of this difference exceeds some small threshold. The fourth input to the control algorithm is the status signal from the interface unit in use. In the normal condition, both phase-locked loops receive their inputs from this interface unit. If the status indication goes off, the phase-locked loops are put into the free-running mode with loop A free running and loop B locked to the output of loop A. This keeps the phase difference between the two loop outputs small even though the system is in the free-running mode. If a slip in loop B is detected when in the free-running mode, appropriate indications are given and both loops will then free run separately since it is not possible to detect which loop is in error.

When in the normal mode, no action is taken when the tracking detector indicates a lack of tracking. However, when the no-track signal is on and a slip is detected, the loop that exhibited the slip is considered to be in trouble and its output is inhibited. The output circuits will automatically switch to the other loop if they are not already in that condition, and a minor alarm is given. If, after this event occurred, the other loop should also indicate a slip, a major alarm is given. However, the output from the second loop is still supplied to the output circuits. Under no condition is the output of both phase-locked loops ever inhibited.

If, in the normal condition, a slip is detected while the tracking detector indicates the two loops are tracking, then the input signal is rejected even though the status indication from the interface unit is good. The timing supply is then put in the free-running mode with loop A free running and loop B locked to the output of loop A.

The nodal timing supply display-and-control panel includes light-emitting diodes imbedded in a block diagram of the system. These diodes indicate the states of various parts of the system.

3.2 Master timing supply

The master timing supply is identical to the nodal timing supply in all respects except for the interface units. The master timing supply is designed to lock on to a standard frequency transmission rather than the framing bit of a DS-1 signal.

A new Bell System Frequency Standard, using primary atomic standards sources, has been installed at Hillsboro, Missouri. The standard frequency transmission from Hillsboro to other offices in the country will be made via a 2.048-MHz pilot tone on L-multiplex facilities. The DDS master timing supply will be installed at St. Louis, which is the nearest DDS hub office to Hillsboro. The standard frequency will reach this office over two separate transmission facilities.

At the master timing supply, new special-purpose interface units are substituted for the normal interface units. These interface units count down the standard frequency to 8 kHz and keep track of the phase of this countdown process with self-checking. Normally, the output derived from the primary route is supplied to the phase-locked loops. If the primary route should have an outage, the input of the phase-locked loops is switched to the secondary route interface unit, in which the phase of the countdown has been previously adjusted to agree with the first interface unit. When the primary route is restored, the loop input is automatically returned to the first interface unit with phase continuity. Should both inputs to the master timing supply fail, the system will free run and the countdown circuits in the interface units will be reset so that upon restoration of either route the phase of the interface unit output agrees with that of the phase-locked-loop output.

3.3 Local (or secondary) timing supply

The local (or secondary) timing supply is a simple and low-cost system intended for use at those DDS local offices that have T1 data multiplexers. It may also be used at small DDS hub offices, which are at the end of the synchronization tree. Its functions are identical to those of a nodal timing supply, the only difference being the use of a simpler phase-locked loop and control. The local timing supply is not able to free run without an excessive slip rate for more than a few seconds. As in the case of the nodal timing supply, fully redundant interface units, phase-locked loops, and output circuits are included. The interface units and the output circuits are identical to those in the nodal timing supply. However, since the local timing supply should not be permitted to free run for more than a few seconds, the interface units are arranged so that the input to the phase-locked loops automatically switches from one interface unit to the other when a failure in one interface unit or line signal is detected.

The fully redundant local timing supply is as reliable as the nodal timing supply. In very small local offices, where cost is the overriding consideration, it is possible to create a nonredundant local timing supply by installing only one interface unit, phase-locked loop, and output circuit. This arrangement, however, is only used where other nonredundant DDS elements are appropriate.

3.3.1 Phase-locked loop

Each local timing supply phase-locked loop includes a temperature-compensated 89A oscillator. This oscillator has a fractional frequency deviation of less than 12 parts per million under the expected extremes of aging and temperature. The 5.12-MHz output of the oscillator is divided down to provide the 8-kHz and 512-kHz outputs required by the output circuits, and to provide other waveforms required by the timing supply.

The phase-locked loop is a simple first-order type with a time constant of 1.04 seconds, so that its frequency response is that of a single-pole low-pass filter with corner frequency at 0.153 Hz. A sawtooth comparator generates a signal proportional to the phase error between the input signal from the interface unit in use and the 8-kHz output signal. The dc value of this signal is used to control the frequency of the oscillator. Only enough filtering is present in the loop to eliminate high-order modulation products from the phase comparator. This filtering has negligible effect on the loop response.

When the natural frequency of the oscillator is off by the maximum allowed value of 12 parts per million, the resultant static phase error of the loop is one-tenth of a cycle. If the input to the phase-locked loop fails, the loop will free run at its natural frequency and the phase will drift by the additional 0.4 cycle required for a cycle slip in 4.16 seconds. If the input is restored before this time, no slip will result.

Whether or not data slips occur depends on the fill of the T1 data multiplexer elastic store.

3.3.2 Monitoring, control, and alarms

As in the case of the nodal timing supply, the local timing supply also includes a display-and-control panel in which light-emitting diodes superimposed on a block diagram display the status of various components of the system. The same phase-metering circuit is also included. If a nonredundant system is installed, a feature built into the alarm logic changes the effect of some troubles that would normally produce a minor alarm to produce a major alarm. A permanent logic input is supplied by the second phase-locked-loop circuit pack. When this circuit pack is missing, the complementary logic level is supplied to the alarm logic. This changes the determination of minor or major alarm conditions and also causes those light-emitting diodes on the display panel that pertain to the second half of the system to be deactivated.

Associated with each phase-locked loop is a monitoring circuit. This monitoring circuit includes a slip detector which performs a function similar to that of the nodal timing supply slip detector and also requires a manual reset. In addition, other malfunctions in the phase-locked loop, in its input signal, or in the monitor itself are detected. In addition to the generation of alarms and indications, the phase-locked-loop monitor causes the phase-locked-loop output to be inhibited in the presence of certain trouble conditions. An end-of-range detector associated with each loop provides an indication when the phase error between the input and output of any loop is greater than one-quarter of a cycle in magnitude.

The only manual controls associated with the local timing supply are control keys that inhibit either of the phase-locked-loop outputs, and the slip-detector reset button. Interlock circuits on the phase-locked-loop outputs assure that under no condition arising from manual intervention or operation of the phase-locked-loop monitors can the output of both phase-locked loops be disabled simultaneously.

IV. CLOCK DISTRIBUTION

The cross-connection scheme for all DDS signals in an office requires that a common 64-kHz bit clock and a common 8-kHz byte clock be supplied to all equipment in the office. The waveforms required by the DDS equipment itself are shown in Fig. 4. The five-eighths duty cycle provides for a maximum tolerance to delays between pieces of DDS equipment and to difference in clock-propagation time from the timing supply to the different pieces of equipment.

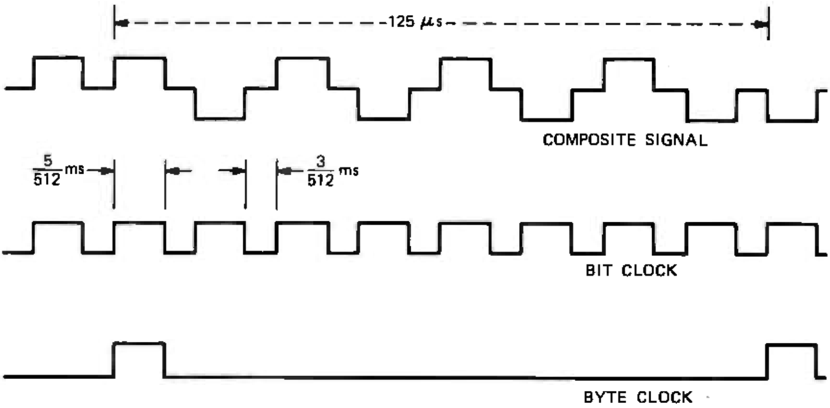

Fig. 4—Output clock waveforms.

The signal distributed by the timing supply to the using bays of DDS equipment contains information for both bit and byte clock combined in a single composite waveform. A composite waveform is transmitted redundantly over two balanced pairs to each using bay. The signal on each pair is a bipolar 64-kHz waveform having a five-eighths duty cycle. Each eighth pulse violates the bipolar rule in that it agrees in polarity with the preceding pulse. The basic waveform, therefore, provides the bit clock information, while the bipolar violation provides the byte clock information. Since the only information is carried in pulse presence and changes of pulse polarity, it is not necessary to keep track of the polarity of the two wires in the transmission line.

Each bay of DDS equipment contains a bay clock, power, and alarms unit, which serves several functions, one of which is to distribute the clock to all equipment in the bay. Two clock line terminators are included in this unit to receive the signals from the timing supply. Each terminator receives signals from both of the redundant pairs. The input composite waveform that is detected first is converted into the two unipolar clock signals required by the DDS equipment. Both terminators will continue to generate the clock signals even when one of the redundant distribution pairs or one of the timing supply output circuits is defective. The appropriate alarms and indications are given if either input signal is defective or if some defect in the circuit operation of the line terminator is detected. A switching circuit connects the outputs from one of the clock line terminators to a number of line drivers that distribute the bit and byte clock to all DDS equipment in the bay. Operation of the switching circuit depends on the presence of detected troubles.

The bit and byte clocks fed to all DDS equipment in an office are in phase, except for slight variations due to differences in propagation time from the timing supply to each bay.

V. ACKNOWLEDGMENTS

The initial synchronization plan was proposed by R. J. Deaton, J. F. Oberst, and R. B. Robrock. The nodal timing supply phase-locked loop was designed by G. Pasternack and J. M. Fallon, and the monitoring-and-alarm circuits by O. Napolitano. The local timing supply phase-locked loop and associated monitoring circuits were designed by R. C. Morris. The development of the interface unit and phase metering circuit was done by T. R. Lawrence and J. Kee. The output circuits and clock line terminators were designed by K. W. Boyd. Analysis of the synchronization network and its possible modes of failure was performed by N. H. Stochel. J. T. Ohlweiler performed the physical design for the nodal timing supply, and J. A. DeSena for for the local timing supply.

REFERENCES

1. P. Benowitz, S. J. Butterfield, M. P. Cichetti, Jr., and T. G. Cross, "Digital Data System: Digital Multiplexers," B.S.T.J., this issue, pp. 893–918.
2. J. W. Pan, "Synchronizing and Multiplexing in a Digital Communications Network," Proc. IEEE, *60*, No. 5 (May 1972), pp. 594–601.
3. C. J. Byrne, "Properties and Design of the Phase-Controlled Oscillator With a Sawtooth Comparator," B.S.T.J., *41*, No. 2 (March 1962), pp. 559–602.

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

Digital Data System:

Digital Multiplexers

By P. BENOWITZ, S. J. BUTTERFIELD, M. P. CICHETTI, JR., and T. G. CROSS

(Manuscript received July 12, 1974)

The two-stage multiplexing hierarchy developed for the Digital Data System is described. Included in this hierarchy are three synchronous time-division multiplexers and a new 64-kb/s cross-connect arrangement that offer both flexibility and simplified administrative procedures. Maintenance for the multiplexers is provided on an in-service monitoring basis and includes automatic switching of a "hot" spare in the event a fault is detected.

I. INTRODUCTION

The Digital Data System employs all digital facilities for both short-haul transmission within a digital serving area and long-haul transmission to interconnect digital serving areas. Except for the local loop which serves individual customers,[1] the T1 carrier system is used exclusively within a digital serving area, while several alternatives are available for the long-haul digital channel. To obtain efficient utilization of these facilities, a two-stage data-multiplexing hierarchy is employed. As illustrated in Fig. 1, each customer's loop is terminated in an office channel unit (OCU) that matches the loop's data rate, e.g., 56, 9.6, 4.8, and 2.4 kb/s. The output of a 56-kb/s OCU feeds directly into a port of the second-stage multiplexer. The other three, which are collectively termed subrate data rates, are gathered in groups of 5, 10, or 20 in the first-stage multiplexer. Each first-stage, or subrate, multiplexer, in turn, feeds a port of the second-stage multiplexer.

Synchronization information is contained within the multiplexed output signal of a first-stage multiplexer. This permits it to operate independently of the second-stage multiplexer. The result is that any port of one second-stage multiplexer may be cross-connected to any spare port of another whether the signal source is a subrate multiplexer or an OCU. To further simplify the cross-connect process, all

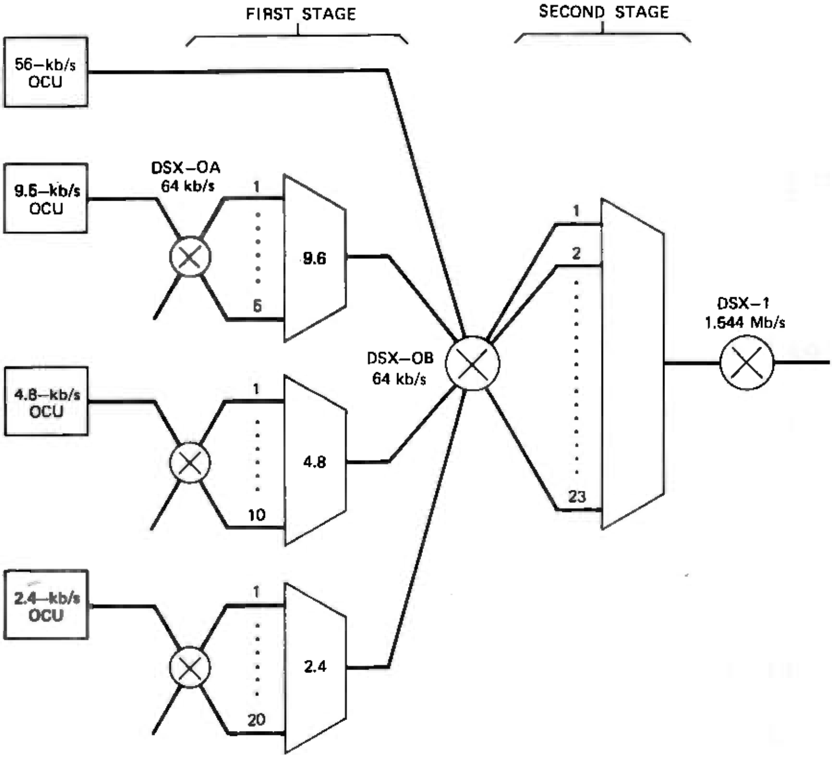

Fig. 1—Two-stage multiplex hierarchy.

signals passing between office channel units and multiplexers do so at 64 kb/s. This arrangement, called a "universal cross-connect" and designated DSX-0, offers wide flexibility and simplified administrative procedures. The signal passing through DSX-0 is called DS-0. The second-stage multiplexer is designed to interface with a T1 carrier system or equivalent digital transmission facilities through a standard cross-connect identified as DSX-1. The signal passing through this cross-connect is designated DS-1.

The failure of a multiplexer, particularly a second-stage multiplexer that may carry as many as 460 individual customer channels, could be catastrophic. To minimize the likelihood of this event, a "hot" spare can be automatically switched into service in the event a malfunction is detected.

II. UNIVERSAL CROSS-CONNECT—DS-0 SIGNAL FORMAT

The Digital Data System employs several different types of transmission equipment that must be interconnected within a central office.

Among these are multiplexers, office channel units (OCUs) for all four data speeds, multipoint junction units (MJUs), and various pieces of testing equipment. To allow a maximum of possible interconnecting arrangements with no tailoring of the equipment to the data rate requires a signal format that is universal to all four customer data rates.

As illustrated in Fig. 2, the format for all intraoffice transmission is organized on an eight-bit byte structure. The eighth bit in each data byte is not made available for the transmission of customer data but is reserved for network control purposes and to assure that the data channel is transparent to the customer's data format. This bit, under control of the OCU, is set to a logical 1 whenever the customer's data terminal is in the transmitting mode and 0 otherwise. Coding the eighth bit in this manner provides the means to transmit control information in the data channel. Any byte in which the eighth bit is a 0 is called a control byte. To meet the requirements of the T1 repeatered line, in particular, items (*i*) and (*ii*) of Section 4.1, it is necessary to make one or more of the other seven bits in a control byte a logical 1. Bit position 1, as described in Section 3.1, is used as a subrate synchronizing channel; therefore, only bits 2 through 7 are available providing a maximum of 63 control codes. Network integrity is pro-

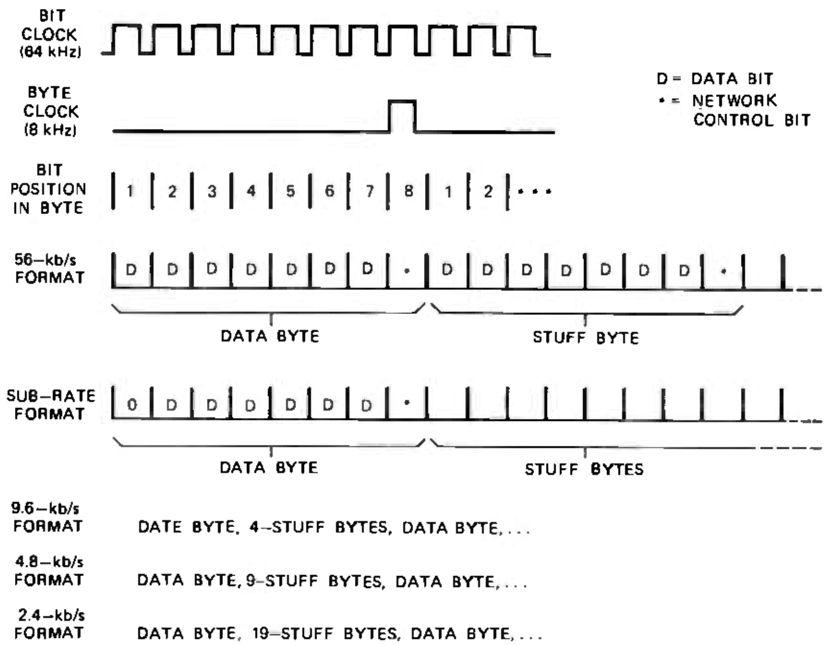

Fig. 2—DS-0 signal format.

tected since the ability to affect the state of the eighth bit is reserved for DDS central office equipment.

The DS-0 signaling rate is 64 kb/s for all data rates; therefore, each byte is 125 μs long. Office channel units for all four data rates translate the customer's data stream into this byte format, second-stage multiplexers accept this byte format for multiplexing onto a T1 line, and all other office equipment operate with the 64-kb/s byte format at their inputs and outputs.

2.1 Data channel capacity

The manner in which the multiplexers operate is to assign all the bits in a given byte to the same data source. Since the eighth bit in each data byte is assigned a network control function, seven bits are available in each byte for the transmission of data. Hence, the maximum data rate is 56 kb/s. Subrate channels are derived by sharing the same 64-kb/s channel among several sources. However, only bits 2 through 7 are assigned to customer data. Since six bits per byte are used, the maximum data capacity per channel is 48 kb/s. If a byte is assigned to a given data source once every five frames, that data source can transmit at 9.6 kb/s. Conversely, one 64-kb/s channel can be used to transmit signals from 5 data sources, each operating at 9.6 kb/s, or from 10 sources, each operating at 4.8 kb/s, or from 20 data sources, each operating at 2.4 kb/s. To distinguish among the 5, 10, or 20 subrate channels that share the same 64-kb/s channel, a synchronizing pattern of length 5, 10, or 20 is inserted into the first bit position of each byte. This is discussed in detail in Section 3.1

Because the 64-kb/s cross-connect data rate is faster than the customer's data rate, the latter must be increased. This is accomplished by repeating the bytes 5 times in succession for 9.6-kb/s data, 10 times for 4.8-kb/s data, and 20 times for 2.4-kb/s data. The repetition interval for each speed is exactly equal to the duration of six-bit intervals at the customer's data rate. For example, at the 9.6-kb/s data rate, six bit intervals equal 625 μs ($6 \div 9600$ seconds). Five byte intervals at 64 kb/s also equal 625 μs (5×125 μs).

The DS-0 format enables considerable flexibility in the design of the overall DDS network. It allows a subrate multiplexer of a single design to be easily adapted for all three subrate speeds, it enables one basic design for the MJU to be used at all speeds, and it provides a maintenance scheme common to all four data rates.

2.2 Timing signal constraints

The DDS is a synchronous transmission system, i.e., timing information is required with every data stream. To eliminate the need to transmit timing information between pieces of central office equipment or to derive it from the data signal, a single source of 8-kHz and 64-kHz reference timing signals exists in each office. All transmitting circuits transmit data to the DSX-0 and all receiving circuits sample data from the DSX-0 at instants defined by the same office clock. The 64-kHz and 8-kHz reference signals are distributed from a central clock source to all equipment in an office. Data are transmitted into the cross-connect on positive clock transitions and data sampling occurs on negative clock transitions of the 64-kHz clock.

The delay in transmitting the clock from the office timing supply to various pieces of equipment and the delay involved in transmitting data between equipment are both significant fractions of the 64-kHz clock period. The difference in the clock delay between the timing supply and two pieces of equipment may be as large as 3.0 μs, and the data transmission delay may be up to 5.0 μs. In addition, there may be up to a 0.5-μs delay in distributing a clock signal within a bay of equipment. The time interval between the write (or transmit) transition and the read (or sample) clock transition of the 64-kHz clock must be greater than the sum of these maximum time delays, e.g., greater than 9.0 μs. To accomplish this, the 64-kHz clock has a $\frac{5}{8}$ duty cycle, i.e., it consists of a positive-going pulse 9.8 μs in duration every 15.67 μs.

The 8-kHz office clock is distributed to all equipment to define the interval between 8-bit bytes. This clock has a duty cycle of 5/64 (9.8 μs every 125 μs). As illustrated in Fig. 3, the 9.8-μs pulse in the 8-kHz clock coincides with every 8th bit of the 64-kHz clock. The two reference timing signals are distributed in the form of a composite signal described in a companion article on DDS network synchronization.[3]

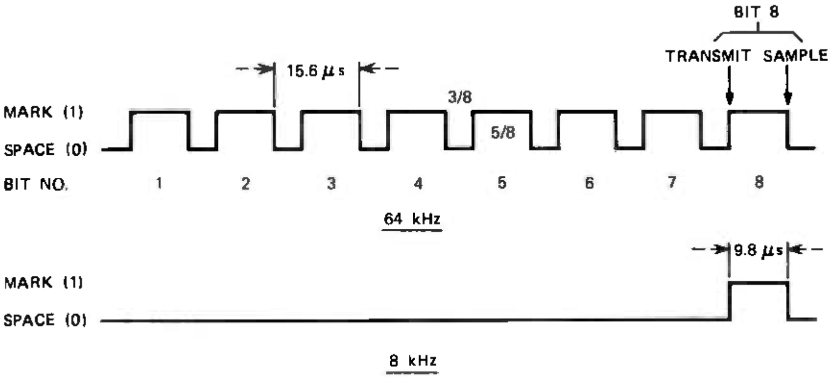

Fig. 3—Reference timing signals for universal cross-connect.

2.3 Electrical characteristics of DS-0 signals

Data are transmitted between DDS central office transmission equipment (e.g., T1DM, OCU, MJU, etc.) in a nonreturn-to-zero bipolar format. Balanced, shielded, twisted-wire pairs are used that are transformer-coupled at each end to reduce the effect of electromagnetically induced noise currents on the operation of active circuits. The need for an inexpensive line driver and terminator plus the requirement for the dc-free signal led to the adoption of a bipolar format. Time delay requirements stated in Section 2.2 limit the length of the cross-connect interconnecting wires to 1500 feet. Attenuation requirements limit the wire-pair loop resistance to 77 ohms.

III. FIRST-STAGE MULTIPLEXER

In the Digital Data System, the first-stage multiplexer is designated the subrate data multiplexer (SRDM). It provides efficient packing of the lower rate channels by combining several subrate data channels operating at 2.4, 4.8, or 9.6 kb/s into a single 64-kb/s channel. Functionally, there are three different forms of a subrate data multiplexer, one for each of the three customer data rates. A 2.4-kb/s SRDM can multiplex up to twenty 2.4-kb/s channels onto a 64-kb/s line; a 4.8-kb/s SRDM can combine up to ten channels of that speed onto a 64-kb/s line, and a 9.6-kb/s SRDM can multiplex up to five such channels onto a 64-kb/s line. One SRDM-multiplexed data stream occupies one 64-kb/s channel of a second-stage multiplexer.

A direct result of the DS-0 format described in Section II is that an SRDM input port can accept a signal whose bit rate is lower than that of the SRDM. For example, a 4.8-kb/s SRDM can have among its inputs any number of 2.4-kb/s channels in addition to 4.8-kb/s channels as long as the total number of channels multiplexed is not more than 10. The 2.4-kb/s channel will be scanned every 10 bytes and will appear twice in the multiplexed stream. At the output of the far-end demultiplexer, the 2.4-kb/s channel will then appear with each byte repeated 20 times. Similarly, a 9.6-kb/s SRDM can multiplex any combination of five subrate signals, regardless of whether their speed is 2.4, 4.8, or 9.6 kb/s. In this way, if a given route must handle, say, only three customers, one at each of the three subrate speeds, then only one 9.6-kb/s SRDM is required.

3.1 Frame structure

The multiplexed frame format of each SRDM consists of 5 (or 10 or 20) bytes, one from each channel. The first bit of each channel's byte is always a logical 0 when received from an OCU or another SRDM. The SRDM multiplexer inserts a sequential pattern into this bit position

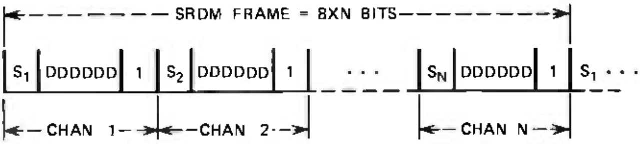

FRAME RATE = $\frac{8000}{N}$ PER SECOND

BIT RATE = 64 kb/s

SRDM DATA RATE	N	S_1	S_2	S_3	S_4	S_5	S_6	S_7	S_8	S_9	S_{10}	S_{11}	S_{12}	S_{13}	S_{14}	S_{15}	S_{16}	S_{17}	S_{18}	S_{19}	S_{20}
9.6 kb/s	5	0	1	1	0	0															
4.8 kb/s	10	0	1	1	0	0	1	0	1	0	0										
2.4 kb/s	20	0	1	1	0	0	1	0	1	0	0	1	1	1	0	0	0	0	1	0	0

Fig. 4—Line format for first-stage data multiplexer.

to form a framing channel. The length of the framing sequence is equal to the maximum number of channels of the SRDM. These sequences are shown in Fig. 4. The five-bit framing pattern used in the 9.6-kb/s SRDM is seen to be a subset of that employed in the 4.8-kb/s SRDM which, in turn, is a subset of that employed in the 2.4-kb/s SRDM. At the demultiplexer output, a 0 is inserted in bit position 1 of each channel.

3.2 Synchronization algorithm

While in sync, the minimum interval to enter the out-of-sync state is approximately 0.5 ms. A searching interval follows which requires three consecutive correct framing bits for a 9.6-kb/s SRDM (four for 4.8 kb/s, five for 2.4 kb/s) for resynchronization. An additional 1.5-ms verifying interval is required, during which correct framing bits are received to ensure that the SRDM is properly in sync.

3.3 Functional block diagram

A functional block diagram of an SRDM is shown in Fig. 5. Data from each of the five channels on a port circuit pack are first multiplexed on that board; data from each of the port circuit packs (if more than one is required) are then combined with the framing pattern on the common logic board and transmitted as a 64-kb/s multiplexed stream. The incoming 64-kb/s multiplexed stream passes through a reframe circuit on the common logic board and is then distributed to the appropriate port circuit packs. The common logic generates an address to each port circuit pack so that the latter can properly distribute its received data to the five channels served by it. When each channel

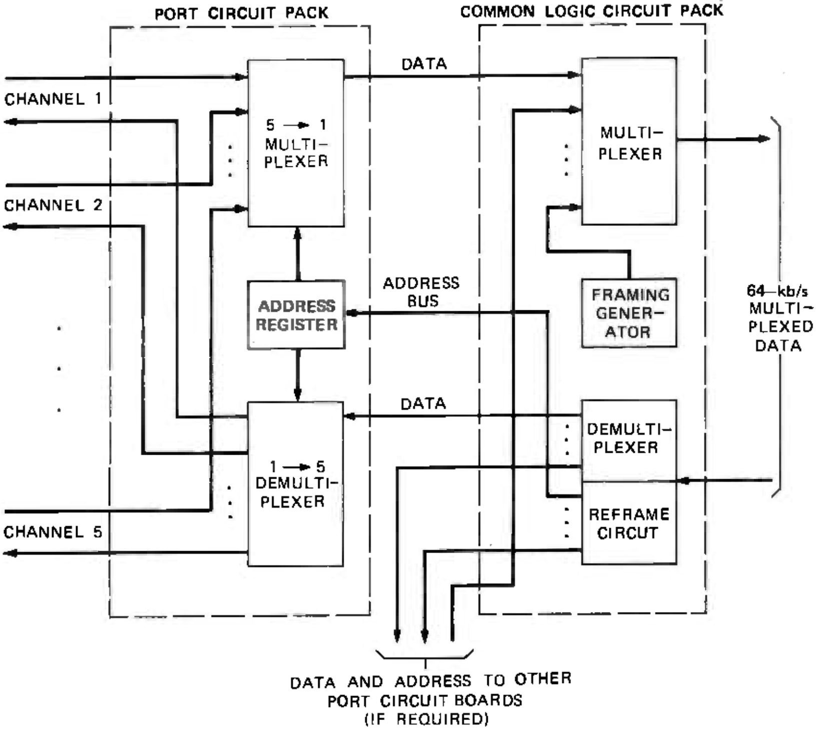

Fig. 5—Subrate data multiplexer, functional block diagram.

circuit receives its byte at the 64-kb/s rate, it is transmitted at the 64-kb/s rate into the cross-connect and is also stored in a recirculating register. The latter provides the means for repeating the byte a total of 5, 10, or 20 times until a new byte is received from the incoming multiplexed stream.

3.4 Integral subrate multiplexer

A special arrangement for a 9.6-kb/s SRDM exists.[1] It is contained on one circuit pack and is called the five-channel integral subrate multiplexer (ISMX). This unit performs the same function as the 9.6-kb/s SRDM but is packaged to mount in an OCU assembly. Up to five channels at any subrate speed can be multiplexed for transmission over a second-stage multiplexer channel without the need for wiring the output of each OCU to an SRDM bay. Unlike the SRDM, the ISMX has no provisions for monitoring, alarms, or automatic replacement.

IV. SECOND-STAGE MULTIPLEXER

Two second-stage multiplexers are available in the Digital Data System. One, designated the T1 data multiplexer (T1DM), is intended for high-usage routes within a digital serving area as well as for operation over long-haul digital facilities. It provides up to 23 data channels and is equipped with extensive performance monitoring, maintenance, and restoration features. The other is designated the T1WB4 data-voice multiplexer (T1WB4). It provides up to 12 data channels and is designed to share a T1 repeatered line with a D3 or D1D voice channel bank over lower usage data routes within a digital service area, or may be used alone if voice sharing is not desired. The T1WB4's more economical design offers somewhat reduced maintenance and restoration features.

4.1 Multiplexer requirements

Basic to the design of the second-stage multiplexers is the need to operate with a T1 repeatered line. Three restrictions are, therefore, immediately evident: (*i*) not more than 15 consecutive 0's may be transmitted,* (*ii*) the average pulse density shall be not less than one pulse out of eight, and (*iii*) the transmission rate must be consistent with the clock recovery circuit found in T1 repeaters, i.e., 1.544 Mb/s. To avoid possible conflict with existing or future T1 carrier maintenance routines, the 193-bit frame employed for voice transmission is retained, as is the coding of the 193rd bit. In addition to these are the requirements of the Digital Data System, namely, to provide data channels that are transparent to all customer data at the four synchronous data rates, to associate with each data channel a signaling channel that can be used to pass network control information, and to operate with a source of timing that is common to all DDS central office equipment.

4.2 DS-1 signal format

The requirements stated in Section 4.1 lead to the multiplexer format illustrated in Fig. 6. The 193-bit frame is divided into 24 eight-bit bytes plus one additional bit designated the F-bit. A byte represents a 64-kb/s channel. When T1DMs are used, all the bytes, except the 24th are data channels. When a T1WB4 is used, the 24 bytes may carry data or digitized voice.

4.2.1 T1DM signal format

The T1DM must operate over both long- and short-haul facilities. In anticipation of possible repeated losses of synchronization because of radio fades or other short-term outages, the synchronizing algorithm is designed primarily to protect against aliasing or false reframing. To accomplish this, six bits of the 24th byte are used in conjunction with

* A zero is the absence of a pulse on the transmission facility.

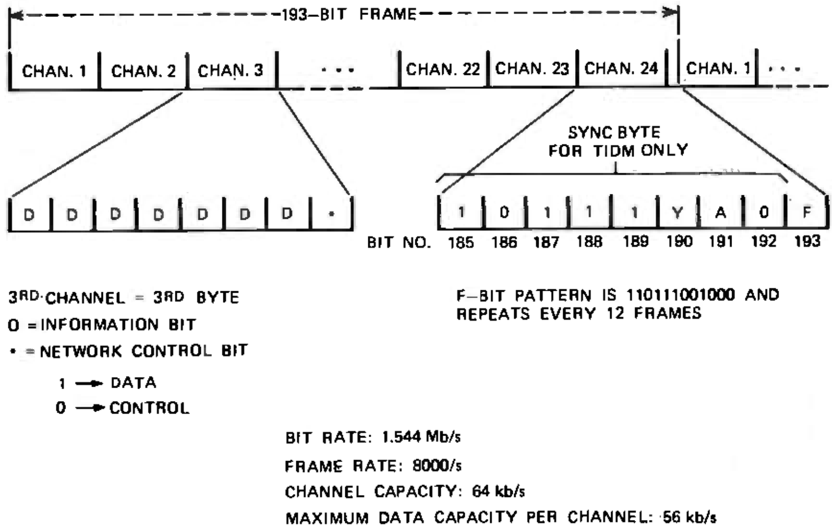

Fig. 6—DS-1 signal format.

the 193rd or F-bit. As illustrated in Fig. 6, the former appear as a fixed pattern while the latter follows the 12-frame pattern employed in the D1D, D2, and D3 channel banks. Bits 190 and 191 are employed for housekeeping purposes by the T1DM. The remaining 23 bytes are assigned as data channels.

4.2.2 T1WB4 signal format

The T1WB4 may share a T1 facility with digital voice channel banks within a digital serving area. To ease the administration of the shared facility as well as maximize its use for voice, the T1WB4 employs only the F-bit for frame synchronization, making all 24 bytes available to carry traffic. While up to 12 of these can be assigned to carry data, there is no restriction as to which bytes carry data and which carry voice.

4.3 T1 data multiplexer

The T1DM combines twenty-three 64-kb/s data channels and a synchronization pattern into a 1.544-Mb/s time-division multiplexed bit stream and converts this to a 50-percent bipolar format[2] (DS-1 signal) for transmission on a T1 repeatered line. It also provides the appropriate demultiplexing functions to convert the received DS-1 signal into twenty-three 64-kb/s channels.

Figure 7 is a simplified block diagram of a T1DM. Each port possesses a four-wire interface with the DSX-0 cross-connect consisting of a bi-

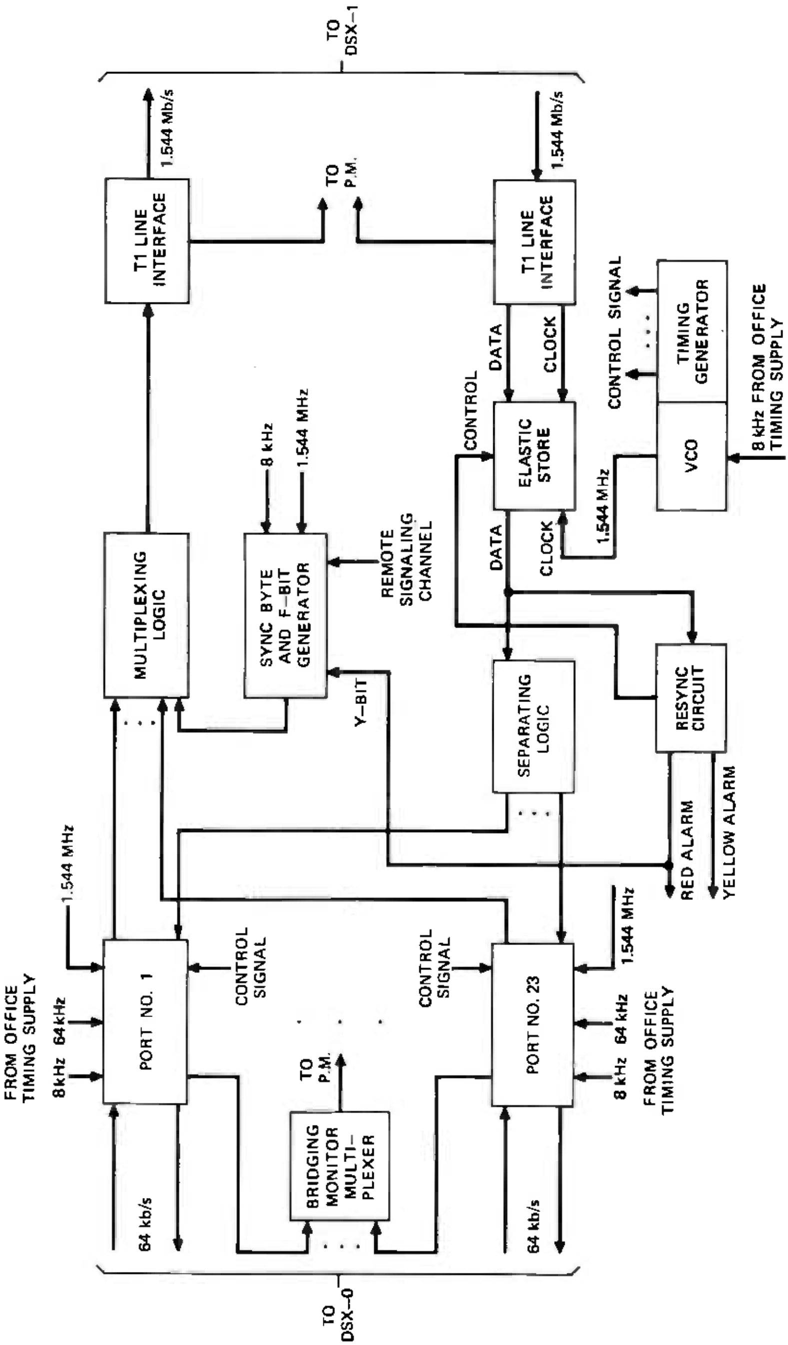

Fig. 7—Block diagram of TIDM.

polar line driver and terminator. The electrical chacteristics of the drivers and terminators are discussed in Section 2.3. Data are clocked into all ports simultaneously under control of the 64-kHz reference signal from the office timing supply. The 8-kHz reference identifies the completion of a byte. Each port contains a pair of 8-bit registers and some control logic. The first register accumulates eight bits of data at 64 kb/s and then serially transfers this byte to the second register under control of an eight-pulse clock burst at 1.544 Mb/s. The byte resides in this register until it is read out at 1.544 Mb/s during the appropriate time slot for each port in the T1 frame. Common circuitry combines these bursts of data, one from each port, and adds on the nine-bit synchronization pattern to complete the 193-bit frame. The transmitting converter converts the logic-level 1.544-Mb/s multiplexed stream into a 50-percent duty cycle return-to-zero bipolar signal for transmission over a T1 line.

After being transformed from a bipolar to a logic level signal in the receiving converter, the incoming DS-1 signal passes through a 256-bit elastic store. The basic purpose of the elastic store is to enable the output 64-kb/s data of the T1DM demultiplexer to be time-aligned with the office 8-kHz and 64-kHz reference signals and to be jitter-free. To accomplish this requires that the phase of the office reference timing signals and the incoming frame be uniquely related. Unless it is provided for, this relation would be arbitrary since the DS-1 signal originates in another central office. One hundred ninety-three bits of the elastic store's capacity is required to align all the data bytes with the office 8-kHz byte reference regardless of the fixed time delay between the sending and receiving central offices. The remaining capacity of 63 bits is available to absorb time variations in the office-to-office delay caused by temperature changes in cable, path changes in radio links owing to atmospheric conditions, etc. The elastic store is basically a 256-bit shift register with a fixed read-in point at the beginning of the register and a readout location that can be varied in one-bit increments anywhere along the register. A 256-state binary up/down counter counts up 1 for every clock pulse from the incoming T1 line interface representing a bit being read into the store, and counts down 1 for every clock pulse of the locally generated 1.544-MHz clock that reads a bit out of the store. The output of this counter provides the address for reading data out of the elastic store. Small differences between the input and local 1.544-MHz clocks will cause the fill of the elastic store to slowly change, but the capacity is sufficiently large to handle all expected variations. If an overflow or underflow (slip) does occur, the fill of the store will immediately be shifted exactly 193 bits in the opposite direction resulting in a deletion or addition of 193 bits in the received DS-1 signal; this is done so the T1DM can maintain synchronization with the incoming DS-1 signal minimizing customer outages in the event of a slip. The 1.544-MHz clock used to read data out of the elastic store is an exact multiple of the office 8-kHz clock and is generated by a voltage-controlled oscillator within the T1DM.

The data that is read out of the elastic store is sent through demultiplexing logic to the port circuits where it is read into one 8-bit register per port at 1.544 Mb/s. The data are then serially transferred to another register at the same speed where it resides until read out into the DSX-0 cross-connect at 64 kb/s. Recall that the input port circuit also required two 8-bit shift registers. Because the input and output 64-kb/s data signals are aligned with the office 8-kHz and 64-kHz references, the multiplexing and demultiplexing halves of the port circuit can each share the same two registers. For each register, one data byte is read in while the other is read out.

4.3.1 Frame synchronization

The output of the elastic store feeds a resynchronization circuit which examines the incoming data stream to locate the seven-bit framing pattern in the 193-bit frame. Six of these are the fixed pattern located in the 24th or sync byte (see Fig. 6). The seventh or F-bit is a 12-bit sequence requiring 12 frames to complete its cycle. While in sync, four out of 12 framing bytes each with one or more of the seven framing bits in error are required to cause the T1DM to assume that synchronization has been lost. The minimum time to enter the out-of-sync state for random signals is 0.5 ms. Once the signal is out of sync, a searching period follows. During this interval, the resynchronization circuit looks for the six framing bits in the sync byte; the F-bit is ignored because its status in previous frames is required to correctly predict its next state. The maximum average time to reframe is about 0.5 ms. Once the six-bit framing pattern has been located, four consecutive frames with correct framing patterns are required to verify that synchronization has been attained. The verifying interval requires 0.5 ms.

During the reframing procedure, the resynchronization circuit controls the delay in the elastic store to force the frame of the incoming data stream to be properly phased with the office byte clock. This is accomplished after the first good framing byte is detected by directing the elastic store to halt its read-out of data until the incoming sync byte is properly aligned with the office 8-kHz clock.

4.3.2 Control codes

The T1DM generates several DDS control words. To comply with the one's density requirement of the T1 repeatered line, the T1DM multiplexer translates any incoming byte that contains only 0's into the pattern 00011000. This is employed within the network to identify a channel as being unassigned, since no OCU can transmit all 0's towards a T1DM. An all-0's channel may have a 1 inserted into the first bit position by a subrate data multiplexer as part of its framing pattern. Although the resulting pattern, 10000000, does not violate the T1 line constraint, trouble isolation would be made more difficult because the channel in question would not have a fixed unassigned code repeating every frame. Therefore, only bit positions 2 through 8 are examined in the T1DM and the pattern 10000000 is translated into 10011000.

The T1DM demultiplexer has two out-of-sync states. If less than 300 ms has elapsed, the incoming bit stream, which may be random data or all 0's, is distributed to the individual ports; if the out-of-sync state persists for 300 ms or more, the T1DM will transmit the control code word 00011010 out of all ports for as long as the out-of-frame condition persists. The latter state is identified locally as a "red alarm." The red alarm condition is transmitted to the far end via the "Y" bit (bit 190 in the frame), where it is received as a "yellow alarm," meaning that the other end's incoming T1 line has probably failed.

4.4 T1WB4 data-voice multiplexer

The T1WB4 has three modes of operation, shown in Fig. 8. The primary mode, Fig. 8a, is to combine digital data signals with PCM voice from a D3 or D1D channel bank for transmission over T1 carrier facilities. In this mode of operation, the T1WB4 acts as a "three-port" multiplexer. It inserts data into unassigned PCM voice channel positions and transmits the combined data-voice signal in a standard format suitable for T1 repeatered lines. The T1WB4 can insert up to twelve 64-kb/s data signals on the T1 line with any or all of the remaining channel positions used for PCM voice. A second mode of operation is the independent data mode shown in Fig. 8b. The third mode of operation is the chained data mode shown in Fig. 8c. This allows data in offices along a route to be added onto the T1 facility without using back-to-back terminals. This mode of operation is made possible by the inherent three-port nature of the T1WB4. Sharing a T1 facility with voice is not permitted in the chained mode.

4.4.1 Functional description

Figure 9 is a block diagram of the T1WB4. The first stage of the transmitter is a receiving converter which accepts a standard 1.544-Mb/s bipolar signal from the channel bank, recovers timing, and converts the PCM voice to logic level signals. The recovered clock is used to drive the counter, the F-bit frame-synchronizing circuits, and the

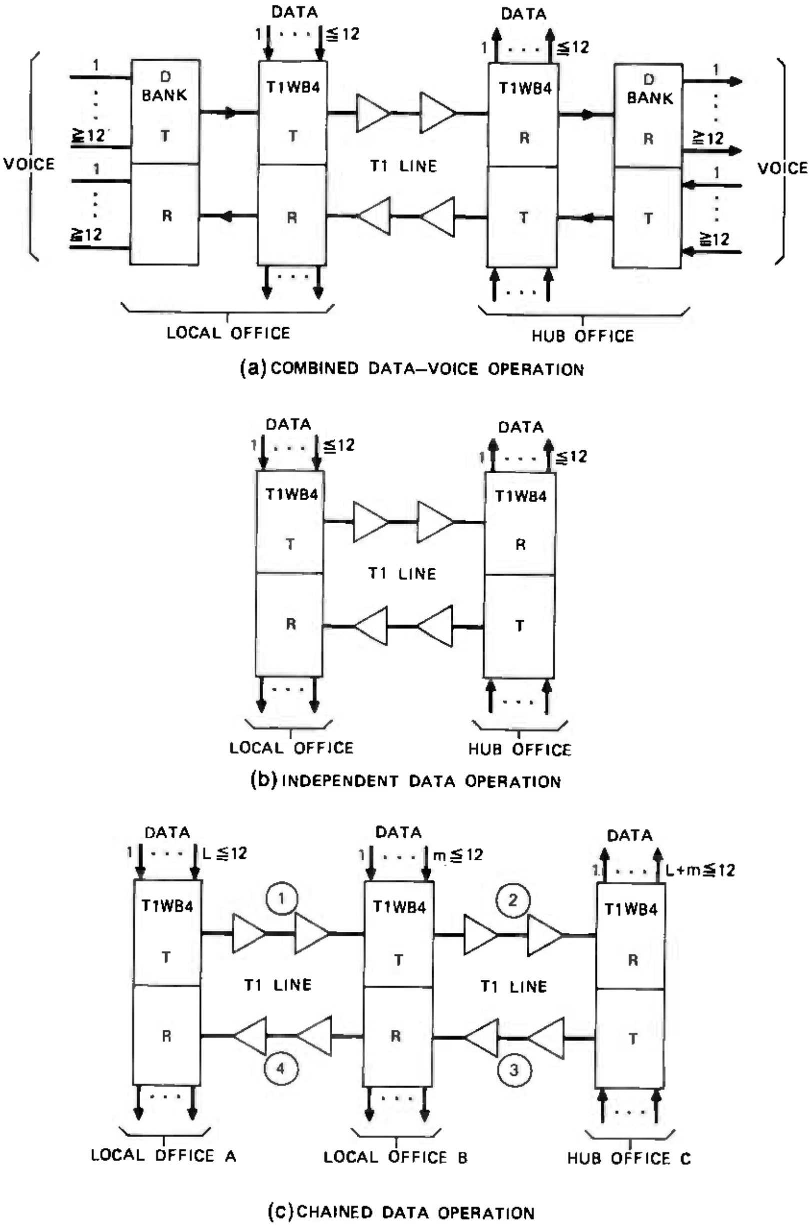

Fig. 8—T1WB4 data-voice multiplexer, modes of operation.

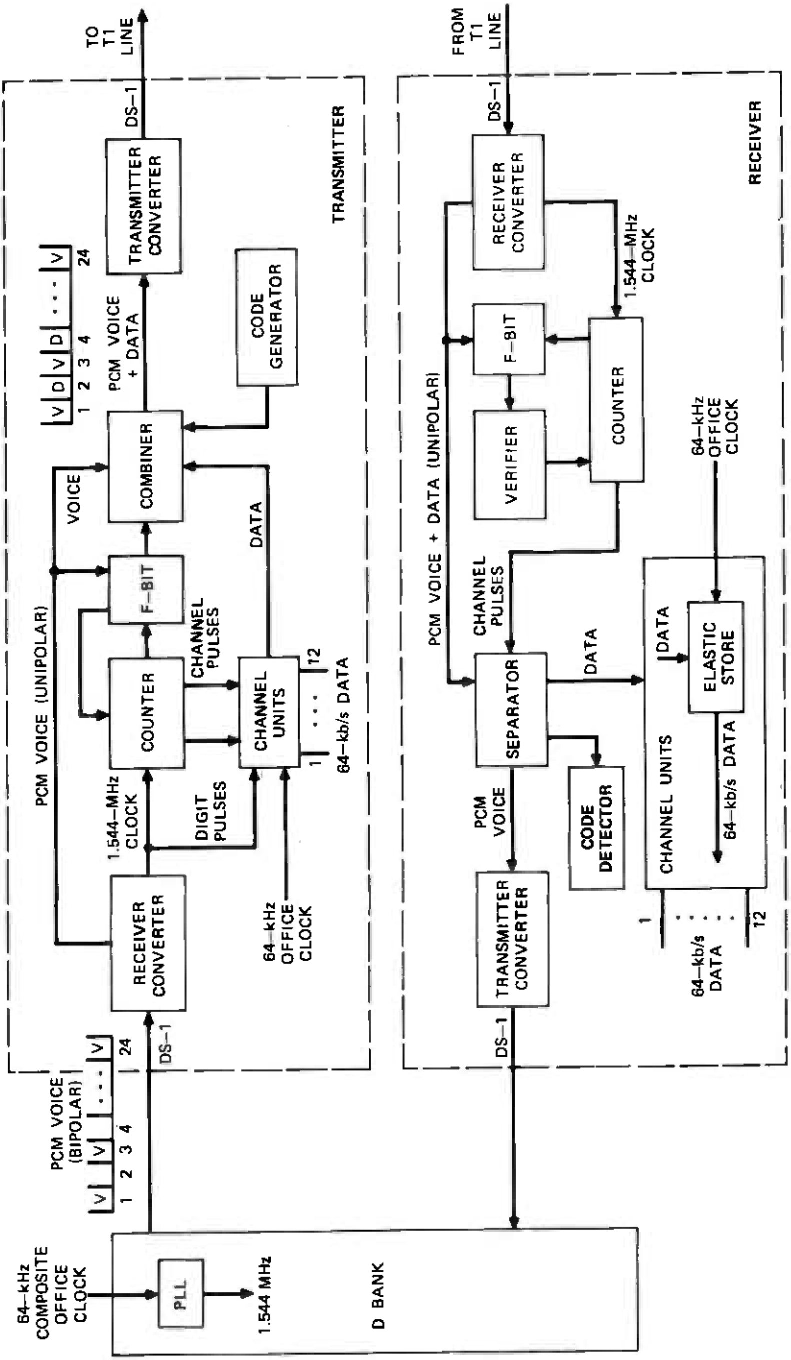

Fig. 9—T1WB4 data-voice multiplexer, simplified functional block diagram.

channel units. Under control of the 64-kHz reference signal from the office timing supply, 64-kb/s data signals are read into all the channel units simultaneously. The data are then sequentially read from each channel unit onto the high-speed bus by the recovered 1.544-MHz clock and combined with the logic-level voice bytes. Following the 24th byte, a single F-bit is generated to complete the 193-bit frame. The multiplexed logic-level signal is then converted via the transmitting converter to the 50-percent duty cycle bipolar format associated with T1 line signals. The receiver functions in a similar manner, but in reverse. The bipolar format of the received DS-1 signal is converted to logic level signals, and the data bytes are separated from the voice bytes and read sequentially into the channel unit receivers. Under control of the 64-kHz reference signal, all 12 data channels are then read out simultaneously at 64 kb/s. The recovered voice bytes are transmitted, as a DS-1 signal, to the D-bank receiver by the T1WB4s transmitting converter.

Per-channel elastic storage is used in the T1WB4 to minimize system start-up cost. The elastic store used in each channel unit transmitter and receiver consists, basically, of two eight-bit shift registers that alternately write in every other frame of data bytes. This approach yields an equivalent elastic storage capability of 96 bits ($\frac{1}{2}$ T1 frame), which is more than adequate to account for line jitter, delay variations in the T1 facility, and the phase offsets resulting from mistuning of phase-locked loops.

4.4.2 Generation and detection of synchronization pattern

Synchronization is achieved by monitoring the F-bit. As noted in Fig. 6, it follows a 12-bit sequence and therefore repeats every 12 frames. The F-bit circuitry is designed to search for the sequence 1, 0, 1, 0, $\cdots$, which occurs every other frame. To guard against the possibility that the T1WB4 transmitter or receiver might falsely frame on a customer data pattern, the verifier circuit examines the remaining six interlaced bits (1, 1, 1, 0, 0, 0) in the 12-bit sequence. A simultaneous match of both sequences is required to satsify the reframing algorithm.

A minimum of three bits out of five checked in the main framing sequence must be in error for the T1WB4 to enter the out-of-sync state. In this state, the sync detector searches matches for the alternating 1, 0, $\cdots$ sequence. Nine successive matches plus verification of the interlace sequence are required for the T1WB4 to enter the in-sync state. The T1WB4 will, on the average, recover sync in about 20 ms.

Two distinct, out-of-sync states are possible depending upon the interval that the T1WB4 is out of sync. If less than 400 ms has elapsed, random data are transmitted from each of the 12 data-channel units. If the T1WB4 is out of sync for 400 ms or more, the out-of-sync code described in Section 4.3.2 is transmitted from all data-channel units.

4.4.3 Zero suppression

Circuitry is provided in the T1WB4 which will replace an all-0's byte with the unassigned channel code (see Section 4.3.2). As with the T1DM, only bit positions 2 through 8 are examined for the all-0's condition in a data byte. This approach cannot be used for voice bytes, because a 1 followed by seven 0's is a legitimate voice code. Hence, the T1WB4 0-suppression circuitry must perform the dual function of examining the last seven bits of each data byte and all eight bits of each voice byte. If eight 0's are detected in a voice byte, the same 0-suppression code word is substituted.

4.4.4 Operation with D-channel banks

In the combined data-voice mode, there is need to maintain data service continuity in the event a D-bank fails. To accomplish this, when the near-end D-bank fails, the associated T1WB4 discards the DS-1 signal from the failed D-bank and establishes its own properly framed signal. This assures good data transmission between T1WB4s. However, this good signal, if passed on to the far-end D-bank receiver would not activate the latter's "red alarm." For this reason, when the near-end T1WB4 transmitter breaks away from its associated D-bank, it inserts a special code (digit 3 missing) in all nondata channels. This code, when detected in the far-end T1WB4 receiver, causes framing to be deleted from the DS-1 signal sent on to the far-end D-bank receiver. This forces the far-end D-bank out of frame and initiates a red alarm. The far-end D-bank transmitter in turn sends a "yellow alarm" code (digit 2 missing in all bytes) toward the failed (near-end) bank. To maintain data service, this yellow alarm code is preempted in the data channels by the far-end T1WB4 transmitter. The near-end T1WB4 receiver detects the yellow alarm code in all the nondata channels and inserts the yellow alarm code in all channels toward the near-end D-bank receiver.

4.4.5 Timing

The T1WB4 timing plan must provide for two functions, namely, the synchronization of the T1WB4 and its associated facilities and the provision, for economic reasons, of an integrated local timing supply to extend the DDS synchronization network to small local offices.

4.4.5.1 Synchronization of T1WB4 and associated facilities. Both the transmitting clock of the D-bank and that of the associated T1WB4 are synchronized to the DDS timing network. In hub offices, this is accomplished by utilizing an external timing source, e.g., the composite clock signal available from the DDS office timing supply.[3] D-banks in hub offices are provided with an interface unit (IU) designed specifically for this application. Loop timing of the D-bank transmitters is employed in local offices. This permits flexible location of the D-banks associated with T1WB4s in local offices, even to the extent of permitting D-banks in remote offices.

In the combined data-voice mode, T1WB4 transmitters in both hub and local offices derive 1.544-Mb/s transmitting clock from the DS-1 signal received from their associated D-banks. In the independent mode of operation, the hub office T1WB4 is synchronized to an 8-kHz signal derived from the composite clock. In the local office, the recovered 1.544-MHz clock from the T1WB4 receiver's input T1 line signal is used to generate an 8-kHz signal. The latter is used to synchronize a 1.544-MHz clock which is then used by the T1WB4 transmitter. Thus, the T1WB4 transmitter is always synchronized, directly or indirectly, to the hub office timing supply.

4.4.5.2 Integrated local timing supply. Section 2.1 points out the need to supply 8-kHz and 64-kHz signals from a common timing source to all DDS equipment. In small local offices, that source is contained within the T1WB4 and is designated the integrated local timing supply (ILTS).

The general structure of the ILTS is shown in Fig. 10. Its input is an 8-kHz clock recovered by the regular T1WB4 receiver from the incoming T1 line signal. A 1.544-MHz oscillator in the ILTS is synchronized to this input. The latter is used as the transmitting clock in

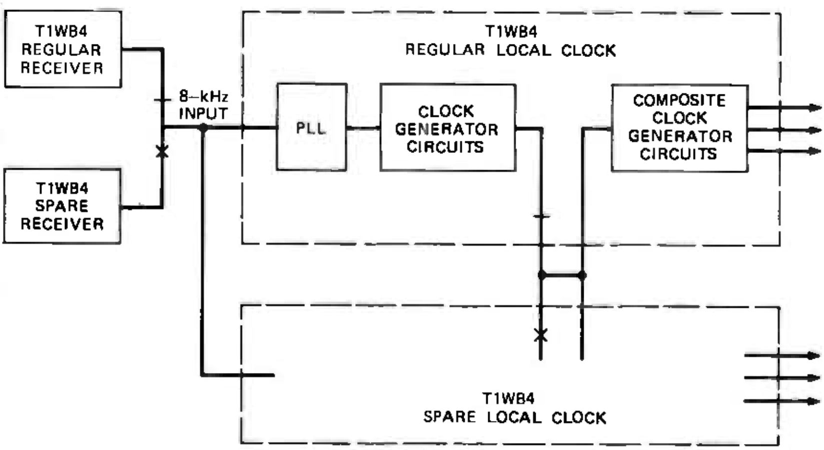

Fig. 10—General structure of integrated local timing supply.

the absence of D-banks and as the input to additional clock-generating circuits in the ILTS. The clock-generating circuits produce 8- and 64-kHz signals that are combined in the composite clock-generating circuit. Three composite clock signals are provided for distribution to DDS equipment associated with the T1WB4.

Two significant features of the ILTS are its ability to bridge short-term outages on the T1 facility (holdover) and its reliability. The ILTS provides the holdover feature by incorporating an oscillator with a moderately accurate free-running frequency and an inexpensive form of frequency memory. This results in a holdover capability of approximately 2 seconds. Reliability of the ILTS is obtained through redundancy.

V. MAINTENANCE

All the data multiplexers (except the ISMX) employ some form of in-service monitoring. The maintenance plan[4] for the Digital Data System is, in part, predicated on the fact that the multiplexers are not only adequately alarmed in the event of a service outage, but include a standby spare that can be used to rapidly restore service.

The data multiplexers take advantage of the digital characteristics of their input and output signals to provide a mechanism for monitoring these signals on an in-service basis. In addition, they provide automatic means to place the spare in service.

5.1 T1 data multiplexer and subrate data multiplexer

The maintenance philosophy for both the T1DM and SRDM is based primarily on the availability of a performance monitor. In the event that a T1DM performance monitor is not available (e.g., taken out of service for maintenance reasons), the T1DM is equipped with alarm and manual test features similar to those found in the D3 channel bank. The SRDM is solely dependent on its performance monitor for testing.

5.1.1 T1DM and SRDM performance monitors

The T1DM performance monitor (T1DM–PM) and the SRDM performance monitor (SPM) provide continuous surveillance of a bay of their respective multiplexers. They detect hardware failures associated with the common circuitry as well as with individual port circuits. Both performance monitors switch into service a spare multiplexer in place of the failed unit, transmit the fault information to the faulty multiplexer for display on a seven-element LED and generate signals to activate the appropriate central office alarms, i.e., minor or major. In the case of the T1DM, transmission failures on the incoming DS-1 signal are also detected and alarmed; however, switching to the spare

T1DM is inhibited. Since both performance monitors are similar in function, only the T1DM–PM is discussed. Where significant differences exist between the T1DM–PM and the SPM, they are noted.

As illustrated in Fig. 11, T1DM–PM access to each T1DMs DS-0 and DS-1 signals is provided through a multiconductor cable. To minimize the number of required conductors, the DS-0 signals from the 23 transmitting and receiving ports are concentrated within each T1DM and appear on two leads. The T1DM–PM sequences through each of the 16 DMs (including the spare) in the same manner. It first checks the multiplexing functions and then the demultiplexing functions of each unit. The multiplexing function is checked by comparing a byte from each 64-kb/s data input port with the corresponding byte, one T1 frame later, in the multiplexed DS-1 signal at the output of the selected multiplexer. The demultiplexing function is checked in a similar fashion with one exception: a byte of 1.544-Mb/s data from the output of the elastic store (rather than the incoming DS-1 signal) is compared, one frame later, with a byte from the corresponding 64-kb/s data output port. Clearly, hardware failures in the data transmission path in either the multiplexer or demultiplexer will result in an unsatisfactory comparison. If the port test fails, the identity of the port under test is stored and the monitor then steps to the next port.

The demultiplexer port test utilizes the 1.544-Mb/s output of the elastic store rather than the incoming DS-1 signal since the delay through the elastic store is unknown and any attempt to compensate for this would overly complicate the performance monitor. Instead, since the elastic store is essentially a long shift register, it is checked by verifying that the sync word can pass through unmutilated. This sync detector circuit is also used to check the transmitted sync word in the multiplexer section of the T1DM.

Aside from the lack of an elastic store test, the SPM's testing algorithm differs from that of the T1DM–PM in another significant area. Before the SPM can begin testing, it must determine the data rate of the multiplexer under test. From this, the SPM can deduce the maximum number of ports to be tested. As described in Section 3.1, three unique framing patterns are used corresponding to a 9.6-, 4.8-, and 2.4-kb/s subrate multiplexer. The SPM determines the data rate by identifying the framing pattern in the incoming multiplexed stream.

Each time that a T1DM performance monitor completes its scan of all the multiplexers, it enters a self-test mode. In this state, the performance monitor is not connected to any multiplexer; hence, the inputs to the monitor self-test circuit should indicate a port failure, a 1.544-MHz clock failure, a transmitted DS-1 sync failure, and the absence of incoming DS-1 signal. If all these events are detected, the performance monitor proceeds to the start of its scan cycle; if not, the

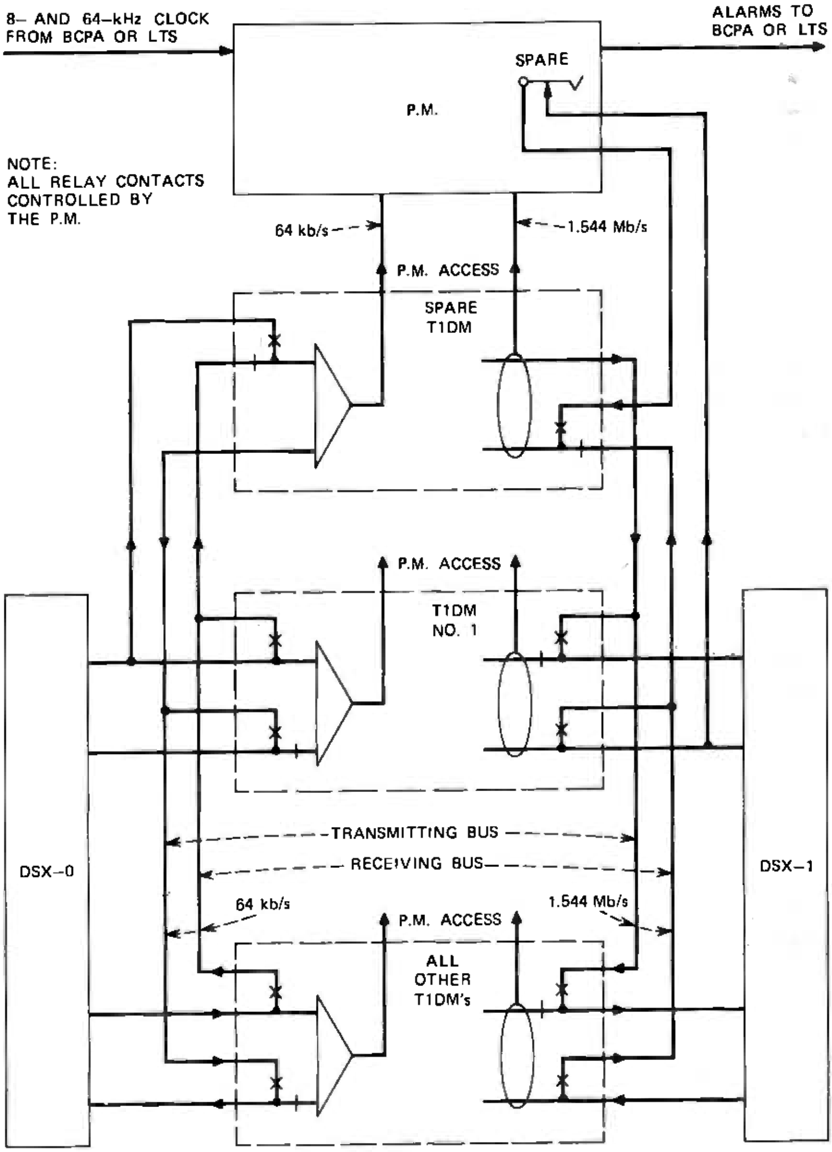

Fig. 11—T1DM performance monitor, functional block diagram.

performance monitor stops and an alarm is initiated. A less sophisticated self-test concept is employed in the SRDM-PM. The scan cycle for a fully equipped T1DM bay (16 multiplexers) is approximately one-half second, while for a SRDM bay (48 multiplexers) it is approximately five seconds.

5.1.2 Standby spare and one-for-n switching

The standby spare is physically and electrically identical to a working multiplexer in every respect. The spare T1DM is checked by the T1DM–PM in the manner described in Section 5.1.1. A source of signals is required by the performance monitor to completely specify the state of a T1DM. To obtain this, as illustrated in Fig. 11, the spare is bridged across a predesignated working multiplexer's DS-0 and DS-1 inputs. Unlike the T1DM, the spare SRDM is checked by a special data sequence generated by the SPM. Also, in the event a subrate multiplexer is found to be faulty, the data rate of the spare is adjusted by the SPM to agree with that of the faulty unit before the spare is switched into service. Thus, the spare is a tested "hot" standby unit. In the event a fault is detected in the spare, a minor alarm is initiated and the appropriate fault symbol is displayed on the spare T1DMs (or SRDMs) LED.

To permit the spare to substitute for any of the 15 other T1DMs, all 16 are bridged onto four buses. Each multiplexer contains a switch module that serves to isolate the multiplexers from the buses. Since balanced transmission pairs must be transferred from a working multiplexer to the spare, economic and reliability considerations dictates the use of relay contacts to accomplish the transfer. Although it is switched out of service automatically, returning a multiplexer to service requires a manual operation. This method offers two advantages: (*i*) it prevents an oscillatory condition from occurring because of marginal operating conditions and (*ii*) it permits coordination with customers where sensitive data services are involved.

As observed in Section 5.1.1, only a hardware failure within a multiplexer will result in a switch to the spare unit. If the spare is already in use, the failure of a second multiplexer will not result in a switch. However, immediately upon returning the first failed multiplexer to service, the spare will be switched in place of the second failed unit.

5.1.3 Alarms and visual aids

In the presence of a performance monitor, all single multiplexer failures, spare as well as working, result in a minor office alarm since customer service is not affected. Both the minor and major alarms result in bay and aisle pilot lamps being lighted. An alarm-cutoff switch provides means to silence both the minor and major audible alarms. If a multiplexer is faulty, its seven-element LED display is energized by the performance monitor. The character displayed identifies which circuit pack or packs are the most likely culprits. A complete list of these characters for the T1DM appears in Fig. 12. A

ALPHANUMERIC CHARACTER	INTERPRETATION	BAY ALARM*	
		YES	NO
123456	SINGLE PORT FAILURE ON CP DESIGNATED	X	
A	AMBIGUOUS CONDITION †		X
C	FAILURE IN CERTAIN COMMON CIRCUITS	X	
E	FAILURE IN ELASTIC STORE	X	
F	FAILURE IN MULTIPLEXER FRAME GENERATION CIRCUIT	X	
P	FAILURE IN SYNC. RECOVERY CIRCUIT	X	
J	FAILURE IN 1.544-MHz VCO IN T1DM	X	
U	RED ALARM – TRANSMISSION FAILURE ON INCOMING T1 LINE, NO SWITCHING OCCURS.	X	
o	YELLOW ALARM – TRANSMISSION FAILURE ON REMOTE T1DM'S INCOMING T1 LINE		X
H	MODE SWITCH IN NO ALM POSITION. IF T1DM PM DETECTS A TROUBLE CONDITION, APPROPRIATE CHARACTER IS DISPLAYED – NO ALARM GIVEN, NO SWITCHING OCCURS.		X
8‡	T1DM OPERATIONAL BUT NOT RETURNED TO SERVICE	X	
BLANK	NORMAL OPERATION		X

*SINGLE T1DM FAILURE RESULTS IN MINOR ALARM AND SWITCH TO SPARE; TWO OR MORE DETECTED FAILURES RESULT IN MAJOR ALARM. IF SPARE FAILS, A MINOR ALARM OCCURS, BUT NO SWITCHING WILL OCCUR.

†T1DM IS ABLE TO RECOVER FRAME SYNC, BUT PM CANNOT.

‡FLASHING

Fig. 12—T1DM alphanumeric displays when performance monitor is used.

similar set is used with the SRDM. If the performance monitor detects a failure internal to itself, it energizes its own displays.

5.1.4 Reliability considerations

The bridging method used by the T1DM–PM ensures the detection of failure of any component (transistor, integrated circuit resistor, etc.). This includes those components that are required to perform the bridging function itself as well as those in the actual data stream. Thus, a multiplexer may fail and a switch to the spare occur because of the presence of monitoring equipment within the T1DM. Also, a failure within the T1DM can cause an unnecessary switch to the spare T1DM. Both these failure modes produce momentary interruptions in service.

A more serious condition is that of service loss resulting from faulty performance monitor operation. While estimates based on *single component* failures indicate that the likelihood is slight that any of these events could occur, they will be detected and an office alarm will be initiated if they should occur.

5.2 T1WB4 data-voice multiplexer

Consistent with its low first-cost design goal, the maintenance arrangement for the T1WB4 is significantly different from that of the T1DM. The T1WB4 provides 1:1 automatic protection switching of the common equipment. No protection of the individual channel unit circuit packs is provided.

To determine when the transmitting or receiving common-equipment circuit packs fail, monitoring circuits are provided at the output of each unit, both regular and spare. This circuitry detects both loss of frame and loss of signal at the T1WB4s DS-1 outputs. By making a comparison of the quality of the input signals to the multiplexer, as determined by the basic framing detection circuitry internal to the common equipment, and that of the output signals, as measured by the monitors, it can be determined when a common unit has failed. For example, if the framing detectors in both the regular and spare receivers indicate a valid incoming signal framing pattern and the monitor at the output of the regular, but not the spare, common receiver indicates a bad output signal framing pattern, the regular common-receiver circuit pack is determined to have failed. This two-out-of-three check using loss of frame/signal indicators provides a good measure of common-equipment operation. If a common circuit pack fails for more than 200 ms, it is automatically spared. Once the failed circuit pack is repaired, it takes about 1 second for the multiplexer to automatically switch back. The delay times are used to guard against over-responding to transient or intermittent conditions.

5.2.1 Automatic breakaway and restoral

The decision to provide 1:1 sparing of the common equipment as opposed to 1:n sparing was based on two factors. First, in the smaller local offices where T1WB4s would principally be used, there would most likely be only a few multiplexers present. A second factor is the need to provide automatic breakaway when D-banks fail. The monitoring scheme previously described makes it straightforward to determine when the incoming signal from a D-bank is bad. If the incoming signal to the T1WB4 transmitter fails because of loss of pulses or loss of framing for 400 ms or more, the T1WB4 will automatically break away and switch to the independent data mode of operation to maintain data service. Breakaway will also occur if the D-bank is determined, by circuitry in the T1WB4, to be asynchronous to DDS timing because of, for example, a failure of the phase lock-loop in the D-bank. To automatically restore to combined data-voice operation, once the D-bank failure has been cleared, requires that monitoring equipment be applied to the incoming line during the breakaway period. Under normal operation, the spare common equipment is used to provide redundant backup for the regular common equipment. When an incoming failure occurs and independent operation is established, the spare common equipment is used as the monitor to determine the status of the incoming signal from the D-bank. This double duty for the spare common equipment results in an inexpensive automatic restoral capability with little if any degradation in system availability since double failures are required before service is affected.

VI. SUMMARY

The Digital Data System provides customer data channels operating synchronously at 2.4, 4.8, 9.6, and 56 kb/s. A two-stage multiplexing hierarchy that obtains efficient utilization of the long-haul 1.544-Mb/s facility has been described. The first stage gathers the three lower data rates in groups of 5, 10, or 20. Each first-stage multiplexer feeds a port of a second-stage multiplexer which is designed to interface with the DSX-1 cross-connect. Up to 460 2.4-kb/s channels can be accommodated in one DS-1 signal. To minimize the likelihood of extensive service interruptions owing to hardware failures, each stage of the multiplexing hierarchy is protected with a standby spare that can be automatically switched into service.

VII. ACKNOWLEDGMENTS

The authors wish to express their gratitude and appreciation to W. J. Mayback and S. B. Pfeiffer for their efforts on the T1WB4, to J. D. Ragucci and R. J. Fretz for their contributions to the T1DM, to G. C. Prins for his work on the T1DM performance monitor, and to T. H. Judd for his work on the SRDM performance monitor. Their untiring efforts and critical comments made possible the successful completion of these units.

REFERENCES

1. E. C. Bender, J. G. Kneuer, and W. J. Lawless, "Digital Data System: Local Distribution System," B.S.T.J., this issue, pp. 919–942.
2. M. R. Aaron, "PCM Transmission in the Exchange Plant," B.S.T.J., *41*, No. 1 (January 1962), p. 99.
3. B. R. Saltzberg and H. M. Zydney, "Digital Data System: Network Synchronization," B.S.T.J., this issue, pp. 879–892.
4. S. M. Fitch and D. L. Rechtenbaugh, "Digital Data System: Testing and Maintenance," B.S.T.J., this issue, pp. 845–860.

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

Digital Data System:

Local Distribution System

By E. C. BENDER, J. G. KNEUER, and W. J. LAWLESS

(Manuscript received July 12, 1974)

The local distribution portion of the Digital Data System is discussed in this paper. Baseband, bipolar transmission over telephone cable pairs is used to extend the digital channel from the serving central office to the customer's location. The performance and requirements of these local channels are presented. The design and performance of the local distribution hardware are examined. Design features include (i) automatic equalization by means of an automatic line-build-out network and (ii) control mode indication by means of bipolar format violations.

I. INTRODUCTION

An important design consideration in planning the Digital Data System was the method of providing local distribution, the interconnection of the customer's terminal equipment to the serving DDS office. Since this interconnection is supplied to each customer, factors such as economy, ease of installation, and reliability were considered.

To minimize the investment required for installation of new telephone plant, a decision was made to use the existing type of subscriber loop plant for DDS local distribution. Thus, the copper pairs from the subscriber's location to the serving central office may be used for DDS as well as for telephone and other services. Since not all local subscriber offices are DDS offices, interoffice trunk cable will also be used as part of the DDS loop.

The local transmission system consists of station equipment, either a data service unit (DSU) or a channel service unit (CSU), to couple the data signals to the cable pairs; two cable pairs, one for each direction of transmission; and central office equipment, an office channel unit (OCU) to convert the digital signals present on the local loop into regenerated bits formed into 8-bit bytes for interconnection to the digital multiplexers.[1]

In addition to transmitting the basic digital signals, means must be provided for transmitting control mode information over the local loop. This mode information may be used by the customer's terminal equipment for control procedures and, as is shown later, is used for internal DDS maintenance and testing originating within the DDS network.

This paper discusses the type of pulse transmission and the automatic equalization used in both the station and office equipment. Performance of the local distribution system is examined in the presence of several degradations. Facility considerations such as cable selection and gauge/length limitations are listed. Maintenance aspects of the local distribution system are discussed. Finally, the equipment is described on a block-diagram level.

II. BASEBAND, BIPOLAR TRANSMISSION

2.1 Bipolar format

The data channel between the station and central office is provided by baseband transmission over the cable pairs. To provide protection against large longitudinal voltages impressed on cable pairs owing to inductive interference from power lines and other sources, it is necessary to isolate electronic circuits from the cable pairs. A reliable and inexpensive means for doing this, which is the method used in DDS, is to isolate the line by means of transformers. The resulting channel has a low-frequency cutoff and therefore requires a signal without dc content. Since unrestricted data patterns can contain a significant dc content, some means must be provided to ensure that the transmitted signal is dc-free. The method used here is to convert the binary data pulse into a bipolar format that removes any dc energy. This method of baseband transmission is used on T1 lines[2] and has proven to be a reliable and inexpensive pulse transmission scheme. Another characteristic of bipolar transmission is the additional information that can be conveyed by means of violations of the bipolar coding rule. As will be shown, this violation procedure is used to transmit the control mode information mentioned earlier.

Encoding binary data into the bipolar format is implemented as follows. A binary 1 is transmitted as a positive or negative pulse with successive pulses alternating in polarity and is called a "bipolar pulse." A binary 0 is transmitted as 0 volts. The "alternating polarity" bipolar rule results in a signal with no net dc component. Figure 1 is a typical bipolar pulse sequence.

2.2 Line driver

A line driver is used in both the station and central office equipment to couple the bipolar signals to the cable pairs. The line driver contains an amplifier, a low-pass filter, and lightning protection circuitry.

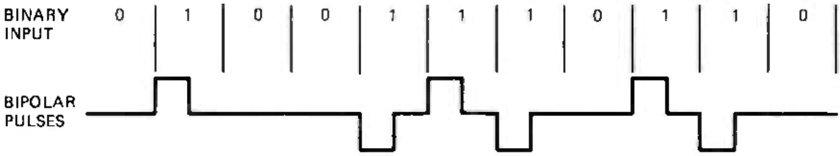

Fig. 1—Bipolar pulse sequence.

2.3 Line receiver

On ordinary paired telephone cable, insertion loss is an increasing function of both frequency and distance. Thus, transmission of high-frequency signals over long distances becomes difficult unless compensatory measures are taken. This is most often done in the form of equalization which counteracts the distributed loss of cable with lumped gain at the receiver.

It is desirable to design a receiver that will compensate the more popular gauges of cable (19, 22, 24, and 26 gauge) over an extensive range of distances. This has been accomplished by using an automatic line-build-out (ALBO) network coupled with fixed equalization within the line receiver. The necessary adjustments are automatic and adaptive. Therefore, installation procedures are simple and the equalization precise, being for the most part independent of the gauge, length, and temperature of the cable pairs.

A simplified block diagram of the line receiver is shown in Fig. 2. The incoming signal passes through the line circuit that contains the line transformer and lightning protection circuits. Noise filtering is provided by a low-pass filter. The ALBO circuit then adds loss to short cable pairs, making all cable pairs appear equivalent in length to match the equalizer. The equalizer adds sufficient gain and frequency compensation to equalize the maximum-length loop. Thus, the combination of cable, ALBO, and equalizer results in a channel with flat loss up to frequencies sufficient to transmit the required bit rate. The output of the equalizer then passes through a three-level slicing circuit

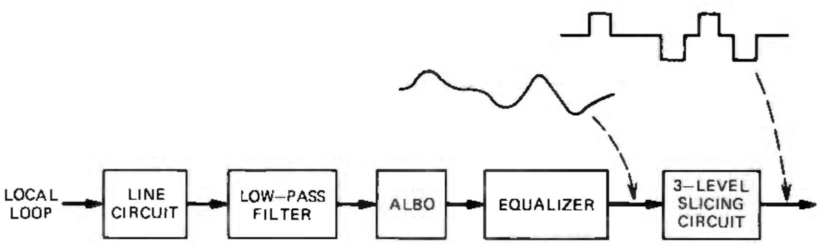

Fig. 2—Line receiver block diagram.

which regenerates the signal in amplitude as shown on Fig. 2. The signal is not sampled by this circuit, so it is not "time-regenerated."

2.4 Operation of ALBO and equalizer

As stated above, the ALBO adds loss to short cable pairs in such a manner as to make every pair appear as a maximum length pair for the corresponding bit rate. Typical loss characteristics of cable pairs can be approximated by a single-pole, low-pass filter in which the flat loss is directly proportional to cable length and the frequency of the pole is inversely proportional to length. The ALBO, shown in simplified form in Fig. 3, consists of an adjustable zero, adjustable flat loss, and a fixed pole. The adjustable zero effectively cancels the equivalent pole of the cable pair. The flat loss combines with the loss of the cable pair such that, with the fixed pole, the combination approximates the loss characteristics of a maximum length loop.

The flat loss and zero location are functions of the variable resistor R, which is realized physically by an FET varistor. Resistor R is inversely controlled by peak detecting circuitry at the output of the equalizer. The peak of the equalized signal has been found to closely track the cable loss at one-half the signaling frequency. Thus, as the associated cable pair becomes longer, the peak signal becomes less, resistor R increases, ALBO loss decreases, and the zero location tends toward the fixed pole. Therefore, on a maximum-length cable pair, the ALBO is essentially transparent, adding neither gain nor loss at any frequency.

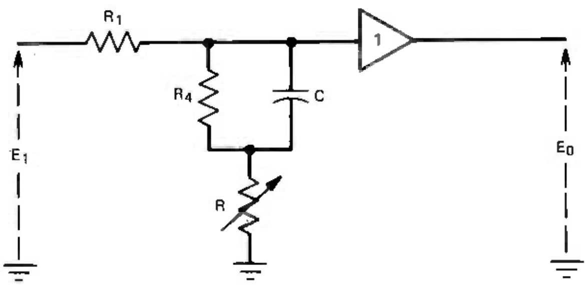

FOR $R \ll R_1$

$$\frac{E_0}{E_1} = \frac{R}{R_1} \cdot \frac{S + 1/R_V C}{S + 1/R_F C},$$

WHERE

$$R_V = \frac{RR_4}{R+R_4}$$

$$R_F = \frac{R_1 R_4}{R_1+R_4}$$

Fig. 3—ALBO simplified diagram.

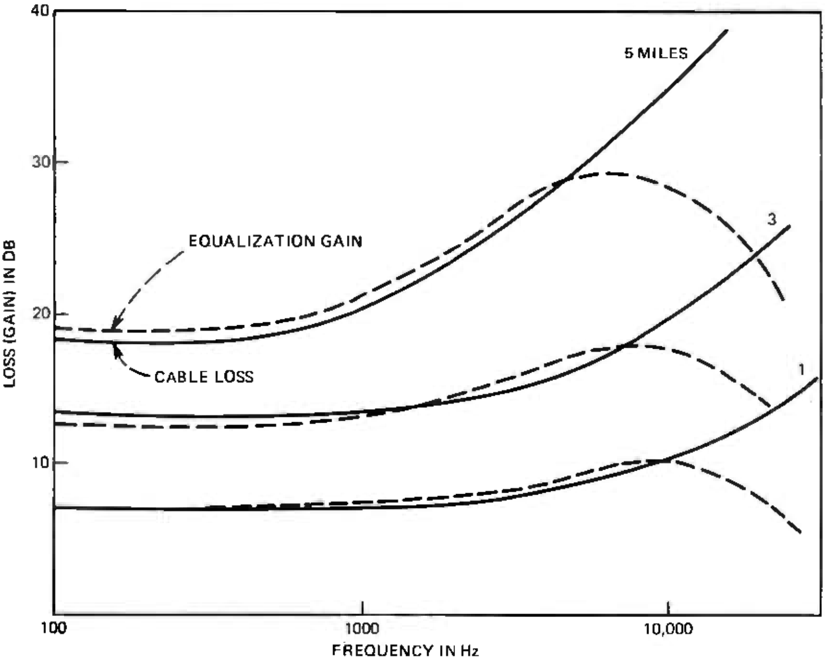

Fig. 4—26-gauge cable loss and equalization for 9.6-kb/s service.

The equalizer is designed to compensate for maximum-length loops. As such, it has fixed gain and a single zero which cancels the effective pole of the ALBO plus cable combination.

Although not explicitly shown in Fig. 2, the line receiver incorporates additional filtering for the suppression of out-of-band noise and crosstalk. To judge how well the ALBO and equalizer match cable characteristics, see Fig. 4 which depicts the insertion loss of various lengths of 26-gauge cable. The inverse of the equalization characteristic for 9.6-kb/s service is also shown, which closely tracks the cable loss over the necessary frequency ranges.

The ultimate measure of performance is eye closure or intersymbol interference. The peak value of intersymbol interference is shown as a function of cable length in Fig. 5 for 9.6-kb/s service on 26-gauge cable. The interference is calculated at the peak of the signal, which is the nominal sampling point, with a value of 100 percent representing a closed eye. The interference is seen to represent only a small fraction of the completely closed eye over a large range in distance.

At long distances, not only does intersymbol interference increase but the possibility exists of unstable operation. For this reason, it has been decided to limit the length of a loop to 31 dB of insertion loss at

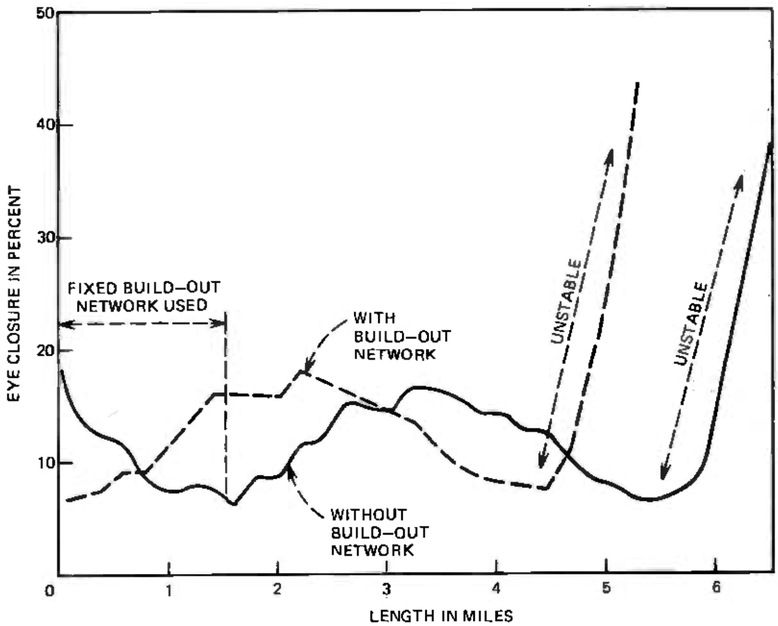

Fig. 5—Eye closure vs length for 9.6-kb/s service on 26-gauge cable.

one-half the signaling frequency. For 9.6-kb/s service on 26-gauge cable, this limit occurs at 4.7 miles, which is well beyond the normal supervision range of most central offices. The additional margin available according to Fig. 5 is used to absorb discontinuities resulting from bridged taps or mixed gauges and temperature effects.

The ALBO is unable to fully compensate for very short distances of cable. As a result, amplifier saturation causes an unacceptable amount of signal distortion. To prevent this situation from occurring, a fixed build-out network internal to the receiver is used for short cable pairs. The additional signal attenuation achieved by this pad allows the ALBO to compensate without saturation and results in the performance indicated in Fig. 5. The build-out network is used whenever the insertion loss of the loop at one-half the signaling frequency is less than 10 dB. On Fig. 5, this loss corresponds to a length of approximately 1.5 miles.

Although it is not shown, performance at other bit rates on 19-, 22-, 24-, or 26-gauge pairs is similar to that of Fig. 5. In all cases, the maximum allowable length of the loop is 31 dB of insertion loss at half the signaling frequency and the fixed build-out is inserted whenever the loop has less than 10 dB of loss.

2.5 Performance and requirements

The previous section considered performance in terms of intersymbol interference. This section considers the effect of outside interference, particularly background and impulse noise, upon the performance of the loop transmission system.

The long-term background noise objective on Bell System loops is 20 dBrn.[3] The noise, which is assumed to be white gaussian noise, is measured across 600 ohms and through a C-message weighting filter which has an equivalent bandwidth of approximately 2 kHz. When this is referenced to the 135-ohm line receiver, an equivalent objective can be found using the following factors:

90 dB	for conversion from dBrn to dBm
$\frac{2}{1}$	for change in bandwidth from 2 to 1 kHz
$\frac{600}{135}$	for change in termination from 600 to 135 ohms.

Thus, the objective at the line receiver becomes

$$20 - 90 - 10 \cdot \log_{10}\left(\frac{2}{1}\right) + 10 \cdot \log_{10}\left(\frac{600}{135}\right) = -66.5 \text{ dBm per kHz.}$$

The equalization shape used to match the cable loss yields a noise gain 3 dB greater than the peak signal gain. This gain is equivalent to the gain at one-half the signaling frequency, which is a maximum of 31 dB. Therefore, an equivalent noise objective at the receiver decision point is $-66.5 + 31 + 3 = -32.5$ dBm per kHz.

Among the four different services, the 56-kb/s service will have the least tolerance to noise because of its large bandwidth. For this service, the effective bandwidth of the receiver, assuming a maximum-length loop, is 48 kHz. Therefore, for cable pairs meeting the background noise objective, the total noise at the decision point is less than -15.8 dBm. With a decision level of 6 dBm, this yields almost a 22-dB signal-to-noise ratio. The associated probability of error, P, is given by

$$P_e = \frac{3}{4} \cdot \left\{1 - \operatorname{erf}\left(\frac{d}{\sigma\sqrt{2}}\right)\right\},$$

where

$$20 \cdot \log_{10}\left(\frac{d}{\sigma}\right) = 22.$$

The factor of $\frac{3}{4}$ arises from the bipolar nature of transmission. That is, a ± 1 is sent from the transmitter one-half the time on the average. In each occurrence, there is a probability of $\frac{1}{2}$ that a noise peak will enhance rather than degrade the signal. Thus, for one-quarter of the time, no error can occur. Evaluating the expression yields a 4-dB margin with respect to an error rate of 10^{-10}. This margin is ample to ensure that errors caused by background noise will be negligible on cable pairs meeting the 20-dBrnC loop objective. Lower speed services of 2.4, 4.8, and 9.6 kb/s require less bandwidth than 56-kb/s service and will have even greater tolerance to background noise.

Although background noise is expected to be negligible, this cannot be assumed for impulse noise. The performance objective for a 56-kb/s DDS local loop is that 99.95 percent of all 1-second intervals should be error-free.[4] Using a conservative assumption that each "second in error" contains only a single error, an equivalent objective on error rate is 0.9 error per half hour. As shown in the appendix, each noise pulse that exceeds the decision level of the receiver causes an average of $\frac{3}{16}$ error. Thus, to meet the error rate objective, noise pulses above the decision level are required to occur less often than 4.8 per half hour.

As was the case for background noise, the decision level is equivalent to $6 - 34 = -28$ dBm at the receiver input. The equivalent bandwidth of the receiver for impulse noise is, however, greater than that found for background noise. This results from the impulse noise bandwidth being a function of the voltage spectrum, whereas background noise is a function of the power spectrum. For the worst-case 56-kb/s receiver, the impulse bandwidth is 64 kHz. Thus, the requirement on impulse noise is less than 4.8 impulses per half hour exceeding a threshold of -28 dBm in a 64-kHz bandwidth.

Measuring equipment commonly used to measure impulse noise, such as the Western Electric 6F noise measuring set, has an effective bandwidth of 34 kHz. Impulse counts in different bandwidths tend to be related as the square of the ratio of bandwidths. Thus, when measured through a 6F set, the objective becomes 1.4 counts per half hour above 62 dBrn (0 dBm = 90 dBrn). This objective is approximately 10 dB more severe than present objectives for cable pairs assigned to 50-kb/s wideband data service.[5]

Lower rate services require less bandwidth than 56-kb/s service, and the objectives are adjusted accordingly. Since an objective of 1.4 counts per 30 minutes is difficult to measure, 10 times as many counts may be allowed for each 10-dB decrease in threshold.

Crosstalk is another factor that must be considered when evaluating loop performance. Again, since crosstalk loss between cable pairs decreases with frequency, 56-kb/s service is the most critical.

At 56 kHz, there is only a 1-percent chance of within-unit near-end crosstalk loss between cable pairs being less than 72 dB. Also, an upper bound on signal power expected from other services is 6 dBm.

Thus, there would only be a small chance of interference exceeding $6 - 72 = -66$ dBm. The decision level referenced to the cable pair is $6 - 34 = -28$ dBm, yielding a 38-dB signal-to-noise ratio. A 38-dB signal-to-noise ratio is ample to ensure that crosstalk from other services will be a negligible source of impairment.

It is also necessary to consider crosstalk from other synchronous DDS services. In this case, the possibility exists of several DDS services, transmitting a periodic pattern synchronously, all at the same bit rate. The worst case would be for 56-kb/s services transmitting an alternating ± 1 signal (steady mark from the customer) on the line. The peak value of this signal would be 12 dBm with almost all the power occurring at 28 kHz. At 28 kHz, near-end crosstalk loss between pairs exhibits a mean of 97 dB with a standard deviation of 8.6 dB for the assumed normal distribution. The 1-percent worst-crosstalk loss would be 77 dB, resulting in a disturbing voltage of -65 dBm. Assuming a decision level-to-crosstalk ratio of 18 dB, which yields a 10^{-10} error rate for gaussian noise, the overall margin available is $-28 - 18 + 65 = 19$ dB.

Using Monte Carlo techniques for summing log-normally distributed powers,[6] it can be shown that this margin is sufficient to absorb 50 multiple disturbers. Since cable pair units typically contain less than 100 pairs, 50 disturbers is the maximum number to be expected in this four-wire service. Thus, adequate protection exists from intrasystem crosstalk.

The conclusions reached above were verified in a field experiment at Freehold, New Jersey. For eight loops (two of each gauge) examined, the C-message weighted background noise was measured to be less than 0 dBrn. With respect to impulse noise, the worst-case loop was a 24-gauge loop with long-term noise as shown in Fig. 6. Data errors and "seconds-in-error" were also measured concurrent with impulse noise. The error rate was found to be approximately 7×10^{-9}, with 99.97 percent of the 1-second intervals being error-free. Thus, background noise, impulse noise, and the resulting error performance were all within objectives.

2.6 Facility considerations

To ensure successful loop transmission, consideration must be given to the characteristics of presently available loop plant. The maximum allowable loop distance is a function of the particular cable gauge and service rate being considered, as shown in Table I. Mixed gauge loops are permitted with maximum distances determined by linear interpolation within Table I. Temperature considerations yield less than a

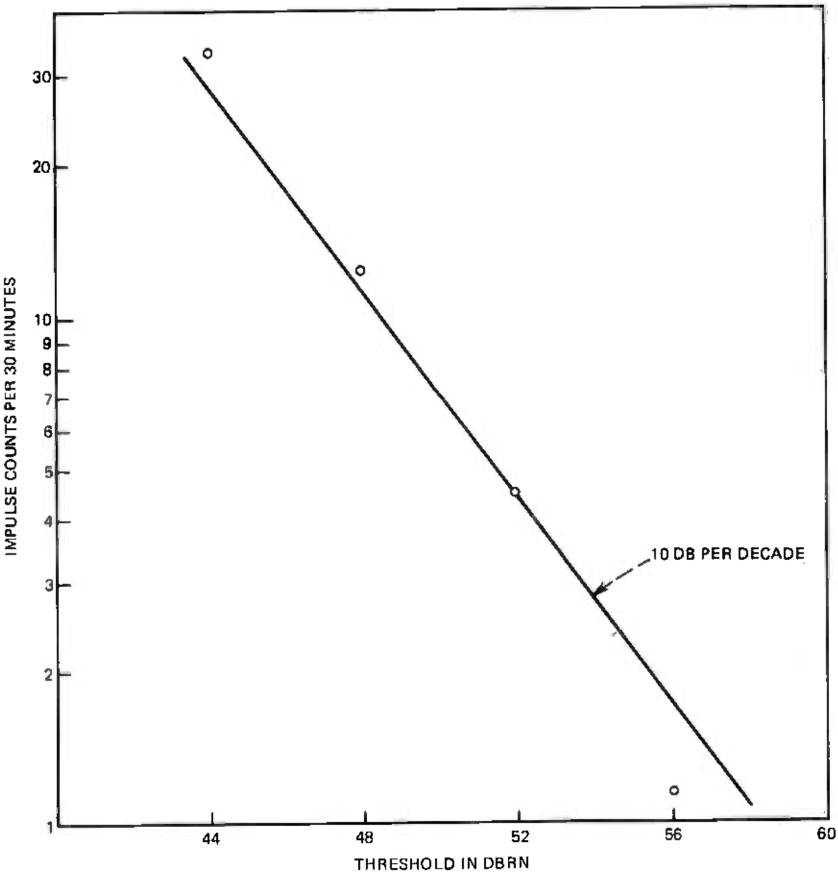

Fig. 6—Impulse noise on worst-case loop.

2-percent reduction in length for every 10°F change in maximum temperature above the nominal 70°F.

Data transmission can be seriously impaired by perturbations in the nominal transmission characteristics of cable pairs caused by loading coils, build-out capacitors, and excessive bridged tap. Therefore, such additions will not be permitted on the loop pairs.

Table I — Maximum distance in miles

	2.4 kb/s	4.8 kb/s	9.6 kb/s	56 kb/s
19 gauge	21.6	16.4	12.7	7.7
22 gauge	13.9	10.7	8.1	4.6
24 gauge	10.6	8.0	6.1	3.3
26 gauge	7.9	6.1	4.7	2.4

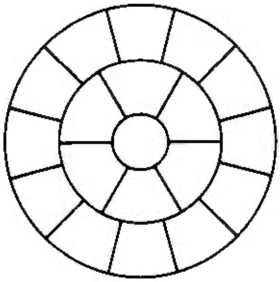

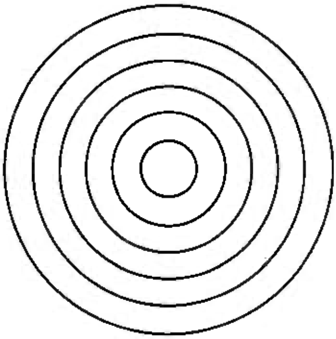

Fig. 7—Cable cross section.

A consideration in providing DDS service is the coordination between DDS and other services that use pairs within the same cable. Certain amounts of segregation are necessary so that undue crosstalk from DDS, particularly at 56 kb/s, is not accumulated in other services. Segregation is achieved by placing susceptible services in adjacent or alternate cable units (splicing groups in the case of layer cable), as in Fig. 7.

2.7 Digital loopbacks

As pointed out in Ref. 7, emphasis has been placed on the maintenance features and testing capabilities of the Digital Data System. Full-time, in-service performance monitoring and protection switching for long-haul and short-haul facilities and multiplex terminals are provided. For the local distribution portion of DDS, alternative strategies present a more attractive cost picture. To maintain the high degree of availability[4] required for the DDS, emphasis has been placed on providing means for very rapid trouble isolation in the event of transmission failure. To accomplish this, digital loopbacks have been incorporated in the local distribution equipment as shown in Fig. 8.

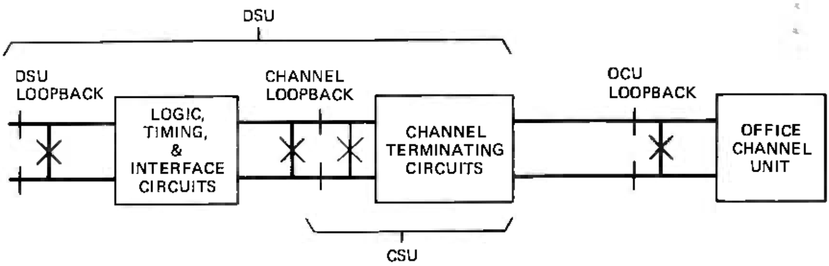

Fig. 8—Digital loopbacks.

In the event of a local distribution trouble, these loopbacks enable rapid isolation of the faulty component so that the proper repair force can be dispatched to quickly restore service. For a channel terminated with a data service unit, three distinct loopbacks are provided. For channels terminated with a channel service unit, two digital loopbacks are provided. Each of these loopbacks can be remotely activated by a distant serving test center (STC). Once activated, digital signals can be transmitted from the STC and looped back to determine the quality of the channel up to the loopback point.

2.8 Network control signals

The private line point-to-point service initially provided in the DDS permits serial binary transmission from end to end without restrictions on the bit patterns. However, for system maintenance within the network and to provide the customer with unmistakable means of indicating and interpreting channel status, additional transmission capacity has been uniquely dedicated to network control information. To protect the control features from accidental or deliberate misuse, the transmission formats have been kept distinct from those of normal data, and customer access to the reserved capacity is strictly limited by the OCU. Thus, the control indication scheme described below applies, in general, only for transmission toward the station.

As described in Ref. 1, the cross-connect format between OCUs and multiplexing equipment (and at tandem multiplexing points) is based on 8-bit bytes as follows:

F1 D2 D3 D4 D5 D6 D7 C8.

Bit F1 is used either for the subrate multiplexer framing code or for data in the case of 56-kb/s service (in which case, the notation D1 will apply in what follows). Bits D2 through D7 are used for data in all services. Bit C8 is dedicated as the network control mode identifier. When bit C8 is a 1, bits D1 (or D2) through D7 are identified as customer data. When bit C8 is a 0, bits D2 through D7 are interpreted

Table II — Network control codes

F1*	D2–D4	D5–D7	C8	Interpretation
1	1 1 1	1 1 1	0	Idle code
0	0 1 0	1 1 0	0	DSU loopback
0	0 1 0	1 0 1	0	OCU loopback
0	0 1 0	1 0 0	0	CHANNEL loopback
0	0 0 1	1 1 0	0	Test code
0	0 0 1	1 0 1	0	MUX out of sync
0	0 0 1	1 0 0	0	Unassigned MUX channel

* Bit column F1 assignments are shown for 56-kb/s operation. For substrate channels (2.4, 4.8, and 9.6 kb/s), this position is dedicated to the subrate multiplexer framing code independent of the content of the remainder of the byte.

as network control information. The various network control codes are listed in Table II. The entries in Table II are grouped according to the content of bits D2 through D4. The first, idle code, stands alone. It indicates no transmission in the customer data stream and is accessible to the user for such purposes as acknowledgment or alerting. The customer is blocked from sending the remaining codes into the network, but they can pass from the network into the local channel for control and information purposes. The group of three loopback codes controls system-initiated remote testing on the local channel. The final group of three indicates trouble or maintenance activities at higher transmission levels of the network, which make the receiving channel unavailable to the customer. By treating the members of these groups as equivalent for local cable transmission, it becomes possible to use bits D5 through D7 to encode the network control information in a form that is distinct from user data. This is accomplished through bipolar format violations without increasing the pulse rate above the service data rate.

In the OCU, bits D1 (or D2 for subrate channels) through D7 are taken from the intraoffice byte and retimed to the service data rate. When bit C8 is a 1, the normal bipolar rule applies, and the resultant line signal carries no indication of the network byte structure. When bit C8 is a 0, a bipolar violation encoding rule is applied. This format violation scheme for network control indication may be explained in terms of the following notation, following Croissier[8]:

- 0 any 0 level pulse
- *B* any ± 1 level pulse transmitted according to the normal polarity alternation rule (the normalized unity level is written here for simplicity)
- *X* a system-determined pulse that may be either a 0 or a *B*
- *V* a ± 1 level pulse transmitted in violation of the normal alternation rule.

When bit C8 is a 0, bits D1 (or D2) through D4 are still encoded by the normal rule, but bits D5 through D7 are replaced with a bipolar violation sequence, X–0–V. The violation pulse, V, uniquely establishes the network control mode and also identifies the byte alignment for interpretation of the network control bits, D2 through D4. The system-determined pulse, X, is set to force the number of B pulses between violations to be odd. This causes successive violations to alternate in sign, thus limiting dc build-up in the transmitted signal. The 0 pulse serves to block occurrences of the sequence, B–V, which can increase intersymbol interference effects.

The encoding rule is probably best seen from an example. Consider two successive idle bytes generated at the far-end OCU in a 56-kb/s system.

$$\left|\; 1 \;\; \underline{1 \;\; 1 \;\; 1} \;\; \underline{1 \;\; 1 \;\; 1} \;\; 0 \;\right|\; 1 \;\; \underline{1 \;\; 1 \;\; 1} \;\; \underline{1 \;\; 1 \;\; 1} \;\; 0 \;\Big| .$$

The underlining emphasizes the grouping of D2 through D4 and D5 through D7 in the control code treatment. The digits are retimed at the near-end OCU in groups of 7 at the 56-kb/s service rate, the four 1s of D1 through D4 transmitted as Bs, and D5 through D7 overwritten with X–0–V.

$$\left|\; B \;\; \underline{B \;\; B \;\; B} \;\; \underline{X \;\; 0 \;\; V} \;\right|\; B \;\; \underline{B \;\; B \;\; B} \;\; \underline{X \;\; 0 \;\; V} \;\Big| .$$

The signed values of the resulting line signal might then be

$$\left|\; + \;\; \underline{- \;\; + \;\; -} \;\; \underline{\overset{(X\,=\,B)}{+} \;\; 0 \;\; +} \;\right|\; - \;\; \underline{+ \;\; - \;\; +} \;\; \underline{\overset{(X\,=\,B)}{-} \;\; 0 \;\; -} \;\Big| .$$

Note that either the sign of the initial B or the use of the X in the first byte might have been different, since the prior state of the system was not given. However, with the first byte established, the second byte is strictly determined with a negative B following the positive V, and $X = B$ to provide an odd number of Bs between Vs. With the violation detected, the specific network control state is identified by bits D2 through D4, which for the example are recognized as the three 1s of the idle code.

The example given for 56 kb/s applies equally to subrate services, except that for the lower rates the F1 bit does not appear on the local loop. Thus, continuous control codes will repeat on a six-bit instead of a seven-bit basis.

As mentioned above, the OCU does not permit the user to freely indicate all network control states back into the 64-kb/s stream. However, an idle code from the station is translated to the corresponding 64-kb/s byte.

One additional violation code occurs with bit C8 in the 1, or data, state. To prevent drift of the receiver timing recovery and ALBO circuits, it is necessary to suppress long strings of 0 pulses on the local loop. For this reason, whenever data bits D2 through D7 are all 0s, the zero suppression network control code is transmitted on line in the form

$$\underline{0\quad 0\quad 0}\quad \underline{X\quad 0\quad V}.$$

At the station, this is recognized and interpreted as a sequence of six data 0s. The same rule applies from the station toward the OCU, except that any string of six 0s is suppressed without reference to byte alignment.

III. OFFICE CHANNEL UNIT

3.1 Format conversion

The principal function of the OCU is to perform a format and rate conversion between the local loop signal at the digital service rate and the 64-kb/s intraoffice signal. The former has been described in detail in earlier sections of this paper. The latter has been described in Ref. 1, but is briefly summarized as follows.

For 56-kb/s service, seven local-loop line bits plus the network control identifier, C8, are retimed to produce a 64-kb/s signal directly. For subrate speeds, six local-loop line bits are grouped between the F1 and C8 bits to produce one 64-kb/s byte. But the local-loop transmission time for six subrate line bits exceeds the 125-μs byte interval at 64 kb/s, so each byte is repeated an appropriate number of times to fill out the additional 64-kb/s byte intervals. For 9.6, 4.8, and 2.4 kb/s, the number of repetitions are 5, 10, and 20, respectively.

The intraoffice signal is transmitted over limited-length balanced pairs through coupling transformers. The signal format is bipolar but without violations. Line driving and terminating circuits are much simpler than those required for the local loops, without closely controlled levels and impedances, shaping filters, or equalizers.

3.2 Block diagram description

Figure 9 is an OCU block diagram. As with all other DDS equipment, operation is synchronous with an office clock source.[9] Clocks are dis-

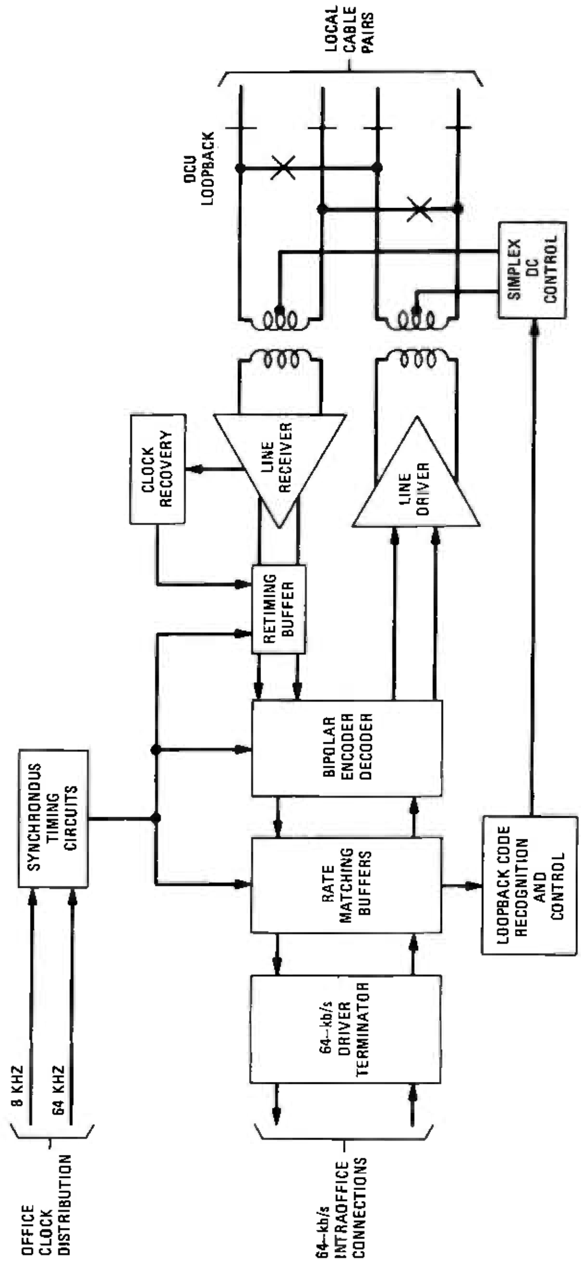

Fig. 9—OCU block diagram.

tributed at 8 and 64 kHz, defining byte and bit timing for the intraoffice 64-kb/s signals. The timing signals for all OCUs are derived by counting down from 1.344 MHz, the least-common multiple of 64 kHz and the four data service rates. A voltage-controlled oscillator driving the countdown chain is adjusted to maintain phase lock with the office clock distribution at the 64-kHz counting level. Byte reference is established to the 8-kHz office clock, and a family of synchronous clock pulses is derived at the 64-kHz and data service rates. Additional pulses are produced once per byte to control data transfer in the rate matching buffers and to establish the bipolar encoder violation sequence. The various paths are abbreviated as single lines in the block diagram.

The rate-matching buffers provide storage to assemble one byte in either direction. On the receiving side, an additional byte of recirculating storage provides the 64-kb/s repetition required for subrate speeds. The bipolar encoder and decoder operate at the service data rate according to the rules given in Section IX.

The line driver and line receiver characteristics have already been described in detail. Essentially identical circuits are used in the OCU, the DSU, and the CSU. The line coupling and protection circuits include the line isolation transformers, lightning protection networks, and the fixed line-build-out pad. Both transformers are center-tapped on the line side, and a small direct current (<20 ma) is simplexed through the local-loop pairs to prevent resistance build-up in unsoldered splices.

Because total delay around the local channel cables is variable, a clock recovery circuit is necessary to establish the data sampling time at the line receiver. A one-bit retiming buffer is used to realign the sampled data with the office-controlled clock of the synchronous timing circuits.

Loopback control codes from the intraoffice connection are detected and distinguished in the recognition circuits as 8-bit bytes. The appropriate loopback state is indicated on the basis of three successive loopback codes received and terminated on the basis of five successive byte intervals without the loopback code indication. This rule applies both at the OCU and the DSU.

In the case of OCU loopback, a relay in the line coupling and protection circuits disconnects the local loop pairs and connects the line driver to the line receiver through both coupling transformers and the fixed build-out pad.

The CHANNEL loopback code is also recognized at the OCU, and a second relay reverses the polarity of the simplexed direct current for detection at the station. During the CHANNEL loopback, violations are suppressed at the bipolar encoder.

In the case of DSU loopback, no special action is taken at the OCU, but the network control code is transmitted with violations for detection at the DSU.

3.3 Multiple OCU arrangement

The block diagram of Fig. 9 indicates the functions of an individual OCU without reference to shared common circuitry. For economy and flexibility, timing circuits and interconnection cabling are organized on the basis of a bay-mounted, two-shelf equipment assembly that can accommodate 20 individual OCUs as plug-in circuit packs. The OCUs for all rates are mechanically interchangeable, with components mounted on two printed circuit boards joined by a common faceplate. The complete circuit pack measures approximately 1.4 × 8 × 10 inches.

The synchronous timing circuits are not included in the OCUs, but are provided in common for each shelf. The 1.344-MHz source oscillator and the basic phase lock circuitry are provided on a single circuit pack at the center of the shelf. A specific rate clock generator on a separate circuit pack is necessary to provide the remaining countdown circuits to generate all clocks necessary for each rate. The specific rate clock generator is inserted to serve five OCU positions on one side of each shelf, which are then dedicated only for OCUs of that service rate.

For 56-kb/s services, the 64-kb/s intra-office driver-terminators for one half-shelf of five OCUs are mounted on a common circuit pack. For subrate services, two arrangements are possible. In hub offices where individual channel cross-connection is required, the driver-terminator circuit pack is used for each group of five OCUs. In an end office where individual channel cross-connection is not provided and efficient multiplexing is not essential, a five-channel integral subrate multiplexer circuit pack may be used in place of the driver-terminators. This circuit pack which is physically interchangeable with the driver-terminator circuit pack enables subrate multiplexing within the OCU shelf.

Power supply arrangements for the OCUs are also multiple. A single dc-to-dc power conversion unit, connected to the office battery supply, provides power for two shelves containing up to 20 OCUs. An OCU power shelf contains two such active power units, plus a working spare, to provide power for four shelves of OCUs. Each power unit contains protective voltage monitoring circuits that control automatic spare transfer and alarm circuits in the power shelf to maintain full service in the event of a single power unit failure.

IV. DATA SERVICE UNIT

The DSU is the more complex of the two customer location units provided in the DDS. It includes clock recovery and logic circuitry to convert between the local loop format and a synchronous binary interface with timing signals provided to the customer terminal. A separate set of control leads in the interface selects the transmission modes (data or idle code) and indicates receiver network control modes. The interface characteristics of the DSU are described in detail in Ref. 4.

A block diagram of the DSU is given in Fig. 10. The line-driver, line-receiver, and clock-recovery circuits are as in the OCU except for the loopback paths. The clock recovery is the overall timing source, since the station is slaved to the network timing. All circuit operation is at the service data rate.

The bipolar encoder and decoder implement the rules given in Section IX. The decoder indicates X–0–V sequences to the network control detection circuitry which identifies the control codes on the basis of the preceding three bits. The DSU idle and out-of-service conditions are indicated after three successive codes are received and cleared after two successive codes are received without the control state indication.

The DSU loopback code is recognized at the DSU according to the frequency rule described under the OCU. The DSU loopback relay breaks the data leads at the customer terminal and provides a loopback path through all DSU line, logic, and interface circuitry.

Polarity reversal of the simplexed direct current from the OCU indicates the CHANNEL loopback command. The reversed current operates a relay through a diode bridge, and a second-stage relay provides the loopback path through the line receiver and line driver without retiming. The OCU suppresses bipolar violations in this mode, as mentioned above, because the untimed regeneration exhibits excessive pulse width distortion in the presence of violations. The use of the line receiver and line driver in the loopback path is necessary, since local channel constraints to ensure desired performance are based on one-way loss characteristics of the local cable. Thus, a round-trip loopback over just the cable pairs could exceed the range limitation. In the DSU, the CHANNEL loopback path is closed directly between the line receiver and line driver to minimize the DSU circuitry involved, so that this test may serve to distinguish between faults in the loop cable and faults in the extensive DSU circuitry bypassed during the test.

The DSU housing[10] measures approximately 12 × 4 × 11 inches. Status lights indicate the presence of required 115-V, 60-Hz power, each of the two loopback conditions, and loss-of-line signal from the OCU. A slide switch permits manual selection of either loopback.

V. CHANNEL SERVICE UNIT

The CSU is the simpler of the two customer location units provided in the DDS. It includes the circuitry necessary to implement the

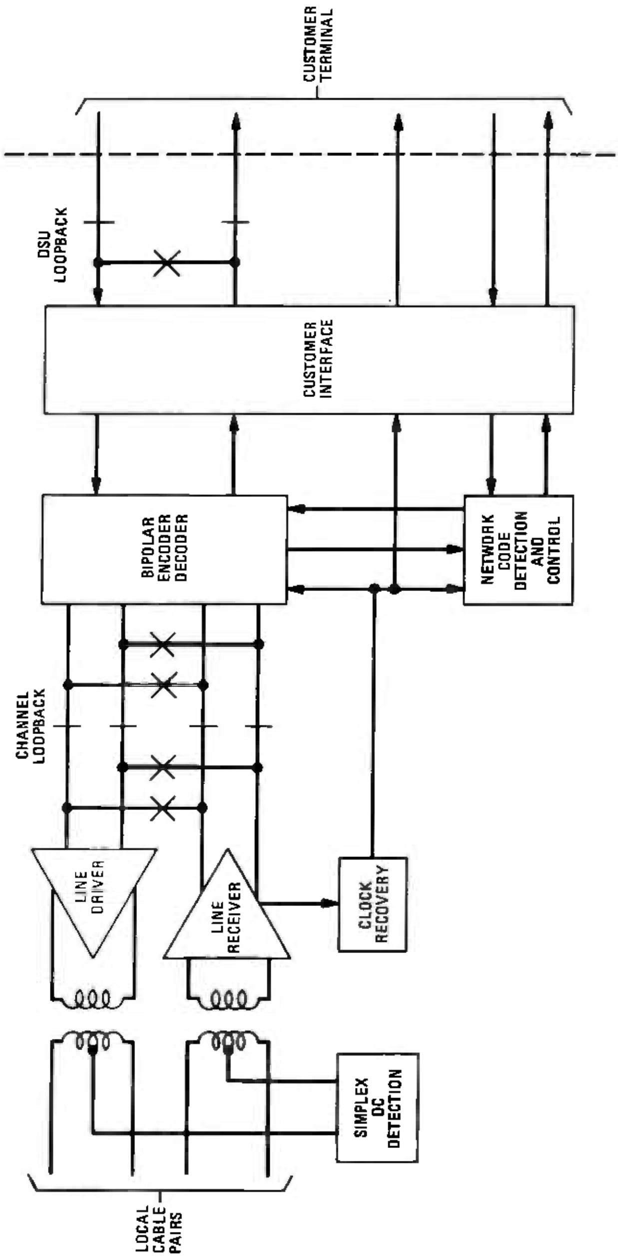

Fig. 10—DSU block diagram.

CHANNEL loopback function required in remote maintenance testing. As noted above, this must encompass the line receiver and line driver. No timing recovery or code conversion circuits are included, and its three-level dc-free interface signals correspond directly to the local channel line signals. Lower cost and simplicity of physical interconnection offer an alternative for the user able to perform his own clock recovery and code conversion functions. The interface characteristics of the CSU are described in detail in Ref. 4.

A block diagram of the CSU is given in Fig. 11. The line driver and line receiver are as in the DSU, including the diode bridge and sensing relay required for CHANNEL loopback.

The digital output of the line receiver is derived from two voltage comparators acting continuously on the filtered and equalized signal. Threshold references for the comparators are set at $\pm\frac{1}{2}$ the nominal peak signal amplitude maintained by the ALBO. The outputs of the comparators directly control a three-level interface driver to provide a replica of the line signal. The slicing action of the comparators yields an amplitude-regenerated signal characterized by distinct transitions between fixed levels. Coupling to the interface is through a balanced isolation transformer. The two outer signal levels are ± 1.4 V, with a pulse duty cycle typically 65 percent of the bit interval but varying with line characteristics and received pulse sequences.

The interface terminator in the transmitting direction also includes a balanced isolation transformer. The customer is required to provide 50-percent duty-cycle pulses to maintain correct pulse shaping in the line driver circuit.

Only the CHANNEL loopback function is implemented in the CSU. Unlike the CHANNEL loopback arrangement in the DSU, the loopback path is closed at the customer interface. Thus, CHANNEL loopback in the CSU serves as a test of both the local cable and all the CSU circuitry.

The CSU housing[10] measures approximately $8 \times 5 \times 3$ inches. Status lights indicate the presence of required 115-V, 60-Hz power and the loopback condition.

VI. SUMMARY

In DDS local distribution between serving central offices and customer locations, the transmission medium is subscriber-loop cable pairs. This paper has described the baseband, bipolar transmission method used to provide the synchronous digital service and its associated control indications over these pairs.

Engineering requirements on the local cable pairs are described in relation to the performance objectives. These requirements include removal of inductive loading, limitations on bridged taps, and a range

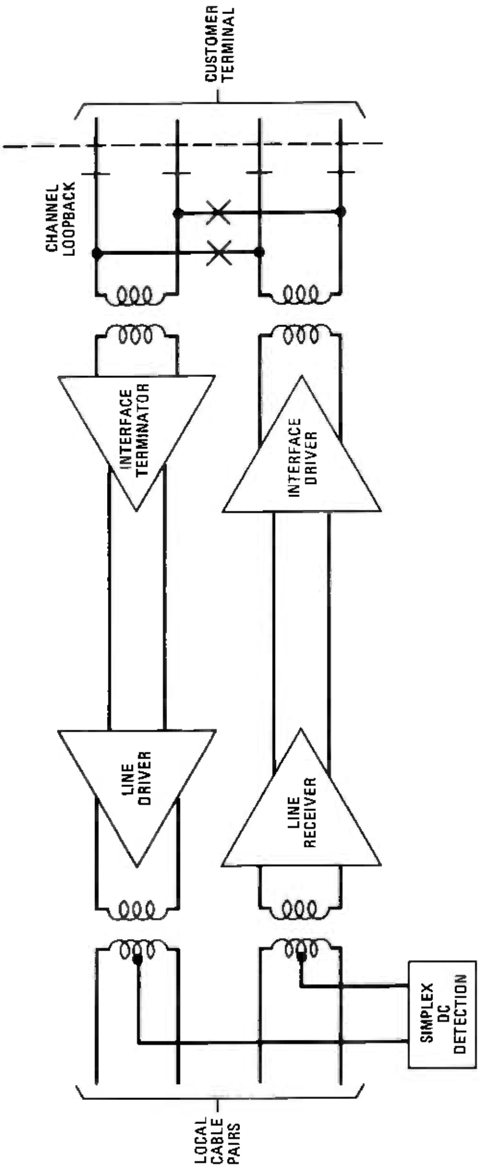

Fig. 11—csu block diagram.

limitation which depends on the data rate and the cable gauge. Reliable attainment of the objectives depends on installation noise tests of the individual loops. Certain sensitive services may experience interference from 56-kb/s DDS loop signals if appropriate separations are not maintained.

Line driving and terminating circuits at the station and the office are similar. All terminators include ALBO networks to compensate for cable length. Station units may include synchronization and logic circuits to provide a standard data interface (data service unit) or may provide only a controlled level bipolar interface (channel service unit). The office channel unit provides rate conversion logic to transform the customer signal to the standard 64-kb/s office format of the DDS. Digitally controlled remote loopback paths are provided at both station and office to permit unaided fault isolation from the serving test center.

VII. ACKNOWLEDGMENTS

The development of the DDS local distribution system was made possible by the combined efforts of many individuals. Specifically, the authors wish to acknowledge the contributions of W. D. Farmer, R. D. Howson, A. L. Pappas, and H. R. Rudy for the circuit design and thorough testing of all the equipment, R. R. Seibel for the programming effort required to provide automated manufacturing tests, R. L. Abbrecht and S. F. Nagy for the construction of early laboratory models, and C. D. Morgan III for his patience and enthusiasm in preparing the manufacturing information.

APPENDIX

Impulse Noise

Several different methods exist that can be used to relate data errors to impulse noise. The one under consideration notes that the receiver appears as a bandpass filter with peak response at the signaling frequency. Thus, impulse noise will tend to appear at the output of the receiver as bursts of a sine wave of frequency equal to the signaling frequency.

Each half-cycle of the sine wave burst has the potential to cause an error and so may be considered a separate noise pulse. From previous experience, it is expected that the long-term number of noise peaks that exceed a given amplitude threshold will increase by a factor of 10 for each 10-dB decrease in threshold. In other words,

$$n(v) = n(v_0)\cdot\left(\frac{v}{v_0}\right)^{-2}, \tag{1}$$

where $n(v)$ is the number of events per unit time greater than v_0 volts. For each half-cycle noise pulse of amplitude v, the probability of an error is proportional to the duration of the pulse above the decision threshold. Note that the maximum duration is one-half the bit interval and assume v_0 represents the threshold,

$$p(e|v) = \frac{\pi/2 - \sin^{-1}(v_0/v)}{\pi}, \tag{2}$$

for

$$v \geqq v_0.$$

Thus, the expected number of errors per noise pulse of peak greater than v_0, $E^*\{N(e)\}$, is given by

$$E^*\{N(e)\} = \int_{v_0}^{\infty} p(e|v)p(v)dv. \tag{3}$$

But $p(v)$ is simply

$$\frac{d}{dv}\left[1 - \frac{n(v)}{n(v_0)}\right].$$

Therefore, evaluating the integral,

$$E^*\{N(e)\} = \tfrac{1}{4}. \tag{4}$$

Note that only $\frac{3}{4}$ of the noise pulses above the decision threshold during the receiver sampling time will cause errors. Thus, the true expected number of errors per pulse is

$$E\{N(e)\} = \tfrac{3}{4}\cdot E^*\{N(e)\} = \tfrac{3}{16}. \tag{5}$$

REFERENCES

1. P. Benowitz, S. J. Butterfield, M. P. Cichetti, Jr., and T. G. Cross, "Digital Data System: Digital Multiplexers," B.S.T.J., this issue, pp. 893–918.
2. M. R. Aaron, "PCM Transmission in the Exchange Plant," B.S.T.J., *41*, No. 1 (January 1962), pp. 99–141.
3. P. A. Gresh, "Physical and Transmission Characteristics of Customer Loop Plant," B.S.T.J., *48*, No. 10 (December 1969), pp. 3337–3385.
4. J. J. Mahoney, Jr., J. J. Mansell, and R. C. Matlack, "Digital Data System: User's View of the Network," B.S.T.J., this issue, pp. 833–844.
5. J. J. Mahoney, "Transmission Plan for General Purpose Wideband Services," IEEE Trans. on Communication Technology, *COM-14*, No. 5 (October 1966), pp. 641–648.
6. I. Nassell, "Some Properties of Power Sums of Truncated Normal Random Variables," B.S.T.J., *46*, No. 9 (November 1967), pp. 2091–2110.
7. S. M. Fitch and D. L. Rechtenbaugh, "Digital Data System: Testing and Maintenance," B.S.T.J., this issue, pp. 845–860.
8. A. Croisier, "Compatible High-Density Bipolar Codes: An Unrestricted Transmission Plan for PCM Carriers," IEEE Trans. Communication Technology, *COM-18*, No. 3 (June 1960), pp. 265–268.
9. B. R. Saltzberg and H. M. Zydney, "Digital Data System: Network Synchronization," B.S.T.J., this issue, pp. 879–892.
10. H. C. Illium, W. B. Lueft, and D. W. Rice, "Digital Data System: Physical Design," B.S.T.J., this issue, pp. 943–964.

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

Digital Data System:

Physical Design

By H. C. ILLIUM, W. B. LUEFT, and D. W. RICE

(Manuscript received July 12, 1974)

The mechanical design of equipment required to build the Digital Data System is described. Economic and technical constraints influencing partitioning, electrical interconnection, and styling are related. An overview is given of the hardware used for terminating loops, multiplexing data streams, and testing for system performance.

I. INTRODUCTION

Communications systems developed to serve a young and growing market and to provide service to a broad geographical area have always presented challenges. Not the least of these challenges is the physical partitioning and packaging of the various subsystems that make up a total service. For the DDS, the task has been to provide private-line data service at several customer bit rates for a changing cross section of users, and to introduce such service into many large cities initially, with a capability of expanding to serve small towns all over the country.

Like most systems faced with interconnecting stations scattered so diversely, the basic DDS layout comprises a hierarchical structure beginning with interconnected hub offices located in large metropolitan areas. Fanning out from these hubs are local offices and the individual subscribers to the service. Equipment in the DDS offices performs the general functions of terminating loops to the customer, multiplexing and demultiplexing data streams, and monitoring and testing the integrity of the system. Station apparatus located at the customer's premises provides the interface between the customer's data terminal and the DDS.

Hub offices are partitioned into four functional sections each of which may be located in separate areas of a central office building or all may be grouped together in one dedicated DDS area.

Local offices have, in general, less equipment and, although they function similarly to hub offices, the integrated nature of their design places all of the equipment in one functioning section. Within certain constraints, the equipment frames in both hub and local offices, though functioning in a given section, may be installed in separate building areas. This has advantages in small installations but for administrative reasons, when floor space permits, one completely dedicated area for DDS equipment, sufficient for all future growth in an office, is more desirable.

All offices may serve customer stations directly; however, customer circuits must be routed to the associated hub office before interconnection to their destination in order that the serving test center located therein may provide maintenance access for rapid evaluation and restoral in the event of trouble on a customer's circuit.

The geographical and functional diversity of DDS equipment posed a challenge in developing packaging schemes that would permit economical arrangement in all locations, keep the amount of new hardware low, and be compatible with the system maintenance plan. This paper discusses the physical implementation of DDS with emphasis on the problems encountered and their solution.

II. TECHNOLOGIES USED

2.1 Components

Circuits for the DDS use conventional discrete components—transistors, diodes, capacitors, etc.—and silicon integrated circuits. The latter, providing either 5 V TTL logic or operational amplifier functions, are used in 16-pin dual-in-line packages (DIPs). Initial production used ceramic packages, with a cutover to plastic DIPs as they became available. One factor that makes possible the high system reliability is the low failure rate of the sealed-junction, beam-lead technology used in Western Electric silicon ICs.

2.2 Printed-wiring boards

The first level of interconnection in the system is provided by double-sided, glass-epoxy, printed-wiring boards. For central office applications, the basic board size is 7.5 by 10 inches which provides for up to 50 connections to other circuitry in the system through gold-plated fingers on the board edge. A maximum of 35 DIPs are mounted on a board, using 0.025 inch as the minimum width for printed conductors and the space between them. More components could be placed on the board if these minima were reduced, but in most cases, it was not economical to do so. For station apparatus, the printed-wiring boards take on special shapes and are described in a later section. In general, however, the boards equipped with integrated circuits were designed with a matrix of wide, printed power and ground paths with filtering capacitors to reduce interference.

2.3 Circuit pack partitioning

In general, there are four conditions that affect the partitioning of DDS circuits into circuit packs:

(*i*) The amount of circuitry on *one* printed-wiring board is compatible with the capability of one 50-pin connector. This results in a circuit pack with one printed-wiring board mating into one connector with a single faceplate.

(*ii*) The amount of circuitry on *two* functionally related printed-wiring boards is compatible with the capability of one 50-pin connector. This results in a circuit pack comprising one "mother board" with a "daughter board" permanently mounted and electrically strapped to it. Only one set of 50-pin connector fingers is provided, those associated with the "mother board," and accordingly, only one faceplate is provided.

(*iii*) The amount of circuitry on *one* printed-wiring board requires more connector capacity than the 50 pins provided. This results in a circuit pack comprising one "mother board" for the components and one or two small, permanently mounted "daughter boards" each providing connector fingers for an additional 50-pin connector. As before, only one faceplate is provided.

(*iv*) The circuitry on *two* printed-wiring boards, though each is independently compatible with the 50-pin connector system, is functionally dependent, thus making it desirable that they be mounted as a unit. This results in a circuit pack with a "mother board" and a "daughter board" each having the 50-pin connector fingers, but both permanently mounted together with only one faceplate.

2.4 Faceplates

When the circuit packs are assembled into a shelf (Fig. 1), the faceplates are the most visible part of the DDS equipment. Thus, appearance and human-factors considerations play an important part in the faceplate design. Switches, test points, indicator lights, and jacks are mounted on the faceplate. Since these are used by maintenance personnel they must be logically grouped and identified. The faceplates are molded of gray polyvinyl chloride, chosen for its appearance, strength, fire retardance, and insulating qualities. Five standard widths

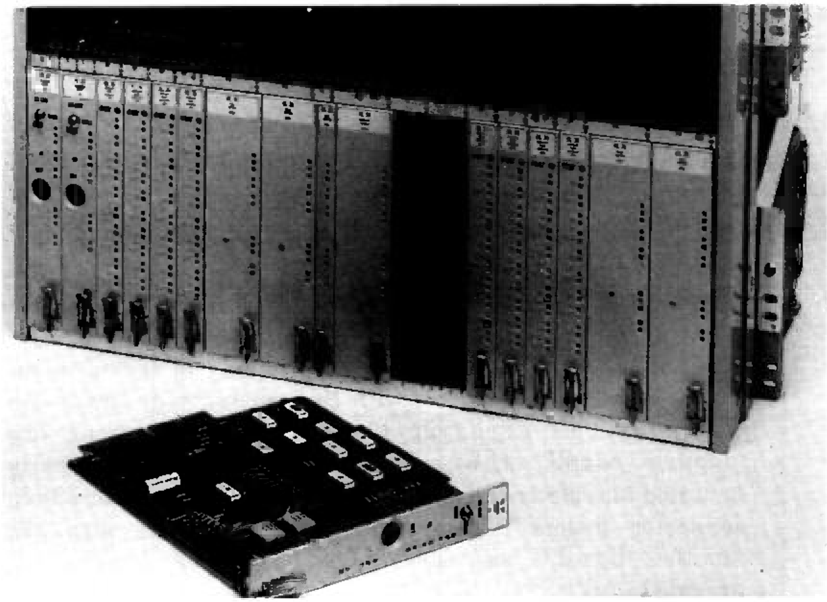

Fig. 1—Typical DDS circuit pack and equipment shelf.

are employed that span the range of component heights used in the system.

2.5 Backplane wiring

The second level of interconnection is between circuit packs, which plug into molded edgeboard connectors. The connector contacts provide early make and late break capability to enable connection to the ground system at those times when power is being connected or disconnected as a circuit pack is being inserted or removed. For each finger contact on the board, the connector provides a terminal post that can be wire wrapped.

On the backplane, the field of these posts at the rear of a hardware shelf, wires are run between connectors to interconnect the boards (Fig. 2). In addition to being a very laborious operation, which contributes significantly to equipment cost, the backplane wiring is a potential source of undesirable interference between conductors. Detailed procedures have been specified for routing wire during manufacture to help ensure that such noise problems do not arise and that a given set of connections are always made in the same manner.

An additional complication on the backplane arises from the distribution of battery and ground. To keep voltage drops low and avoid overloading conductors during accidental short circuits, the use of

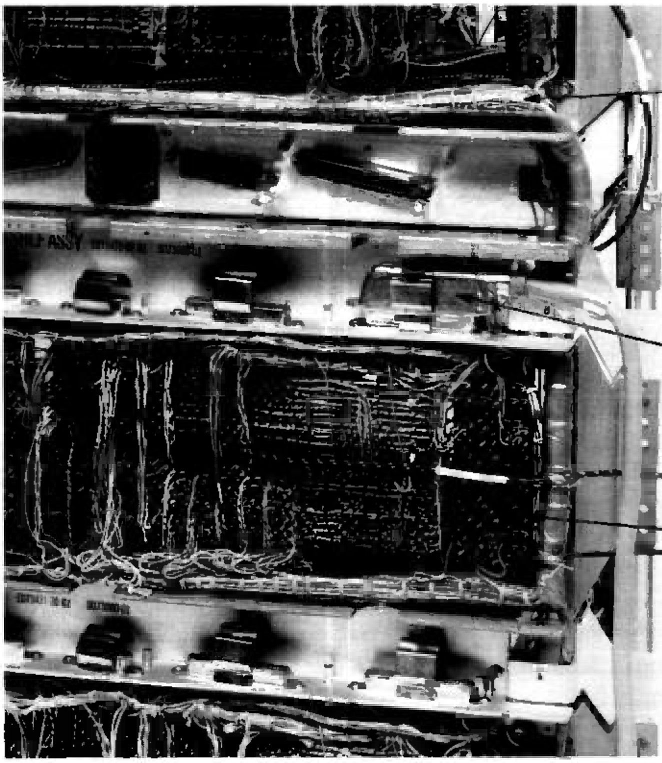

Fig. 2—Backplane wiring on DDS equipment.

large gauge wires was necessary initially in the designs. Since these had to be hand soldered rather than wrapped, adding to the bulk and expense of the backplane wiring, special dual feeds and returns of lighter wire were instituted where possible.

2.6 Intershelf wiring

The third level of interconnection involves wiring between shelves. When the shelves are combined to form a functional, orderable assembly, the wiring is done primarily by factory installed local cables that are terminated directly on the pins of the backplane or on connectors for external interconnections (Fig. 2).

When the wiring involves the interconnection of assemblies to form a subsystem of DDS equipment in one 11-foot 6-inch bay or one or two 7-foot bays, cables with connectors on both ends are used. The connectors mate with those mounted on panels at the rear of the various assemblies. The emphasis here is on simplifying the job for the installer. Rather than to carefully route wires from one location to another, the installer simply secures a connectorized cable of adequate length and plugs it in at both ends.

The following sections describe the application of these technologies to the design challenges presented by several DDS units.

2.7 Equipment shelves

Most of the circuitry in the DDS is provided on plug-in circuit packs, housed in die-cast aluminum shelves like those shown in Fig. 1. Each shelf has 68 printed-wiring-board guides on the upper and lower surface of the shelf base and is arranged for mounting in a 23-inch-wide bay with an overall depth of 12 inches. The shelf itself has no top so that when two shelves are mounted, one above the other, the top guides for a given printed-wiring board are provided on the underside of the shelf above. For single-shelf applications, a die-cast aluminum cover provides the top guides. A maximum of 36 board connectors may be arranged on the rear of the shelf, but the number of guides for the boards allows for extensive flexibility in the choice of circuit-pack widths.

Because of desired circuit-pack partitions and in order to take full advantage of the two-piece characteristic of the shelves, which at minimum vertical separation house circuit packs approximately 5.5 inches high, DDS circuit packs were chosen to be approximately 7.5 inches high. The remaining 2 inches above the printed-board connectors in the rear are then occupied by special brackets for mounting cable connectors and/or terminal strips.

2.8 Power units

Throughout the system, circuits requiring various voltages (+24, +12, +5, −12 or −24 V dc) are used. Central office equipments operate from a primary battery source of either −24 or −48 V. Ferroresonant power converters have been selected and equipment partitions chosen to minimize the number of different power units required.

A major constraint in the DDS is that the failure of a single power unit must result in an alarm with no loss of service on any customer channel. Accordingly, when circuits are totally protected by 1-for-N protection-switching schemes, and power needs are great enough to justify it, power units are utilized on a one-per-circuit basis. In addition, circuitry that is simply redundant is arranged so that each independent half obtains its power from a separate power unit. Circuits with low power requirements, which allow several to operate from one power unit, are arranged so that two such power units, each serving many circuits, are protected by a third "hot" spare. The spare can, without interruption, take over the load of either of the first two power units if an alarmed failure condition occurs.

To meet these central office demands, four basic power units are provided, each available with either −48-V input or −24-V input capability. The four include one with a 50-W output capacity at +5 and −12 V, one with a 100-W output at +5 V, one with a 100-W output at −12, +5, and +12 V and one with a 100-W output at +24 V.

III. CENTRAL OFFICE EQUIPMENT FRAMES

In general, the assemblies, the equipment frames, and the interconnection methods used have resulted in a physical design that provides for easy ordering, short installation intervals, and flexible system configurations. This flexibility includes provisions for using much of the same hardware in both hub office applications and local office applications even though the system requirements and design constraints are quite different for each. Growth capability has been a prime consideration and system rearrangement during growth phases has been facilitated by coordinated connectorization and cabling methods.

The assemblies have been designed and coded so that they may be ordered and installed without regard for the standard system arrangements; however, specific assembly arrangements have been coded which provide bays, engineered as subsystems, that can be ordered partially or completely equipped, installed easily, and interconnected to form the various operating systems required in the overall DDS network. The bays are available in 11-foot 6-inch or 7-foot sizes using unequal flange cable-duct type frames.

In accordance with the subsystem approach, there are essentially nine subsystem configurations. These nine subsystems are:

(*i*) Office-channel-unit arrangements.
(*ii*) Subrate-data-multiplexer arrangements.
(*iii*) T1 data-multiplexer, large-office arrangements.
(*iv*) T1 data-multiplexer, local-office, initial-bay arrangements.
(*v*) T1WB4 data-voice-multiplexer, hub-office arrangements.
(*vi*) T1WB4 data-voice-multiplexer, local-office arrangements.
(*vii*) Nodal or secondary timing-supply arrangements.
(*viii*) Multipoint-junction-unit arrangements.
(*ix*) Serving-test-center arrangements.

Though most of the subsystems are available in both 11-foot 6-inch and 7-foot bay sizes, there are certain differences in the actual quantities of the various circuits included when the subsystem is supplied in one size or the other. For 11-foot 6-inch bay installations, each of the nine subsystems is provided in a single-bay frame. For 7-foot bay installations, they are each provided in one of three ways: by a single-bay frame, by a two-bay arrangement using a double-bay frame, or by a two-bay arrangement using two single-bay frames. The two-bay arrangements using two single-bay frames provide the option of adding the second bay at a later date as growth requires. A typical bay arrangement is shown in Fig. 3, which indicates the standard configuration for office-channel-unit bays.

All DDS assemblies are arranged for front mounting. The mounting holes in the bay framework and on the equipment are arranged so that the equipment may be mounted in vertical increments of $\frac{1}{2}$ inch.

On the rear of each bay framework, with the exception of the cross-connect bays used in the serving test center arrangements, isolated ground busses are provided by insulated vertical rods so that each equipment assembly can be connected to any of three ground systems. The three ground systems are signal ground, battery return ground, and frame ground.

3.1 Office-channel-unit arrangements

The office-channel-unit (OCU) subsystem comprises one bay clock, power, and alarms (BCPA) shelf for distributing power and timing and for accumulating bay alarm signals; three three-shelf OCU and power supply assemblies, three associated two-shelf OCU assemblies; and, when required for specific local-office applications, up to six subrate-data-multiplexer jack-and-connector panels (SM-JCP's). Each of the OCU assemblies can accommodate 20 OCUs and, accordingly, the subsystem has a maximum capacity of 120 customer channels. This subsystem is provided in a 7-foot two-bay arrangement. By deleting one three-shelf OCU and power supply assembly, one two-shelf OCU assembly, and, accordingly, two SM-JCPs, the subsystem is provided in one single 11-foot 6-inch bay with a maximum capacity of 80 customer channels (Fig. 3).

When employed in hub offices, this subsystem is never provided with SM-JCPs since jack access is obtained at the testboard in the same building. When employed in local offices, the SM-JCPs are provided only for OCU assemblies operating at subrate speeds, in conjunction with subrate data multiplexers (located in other bays). An additional attribute of this subsystem is its use, in local offices, of 5- and 10-channel integral subrate multiplexing.

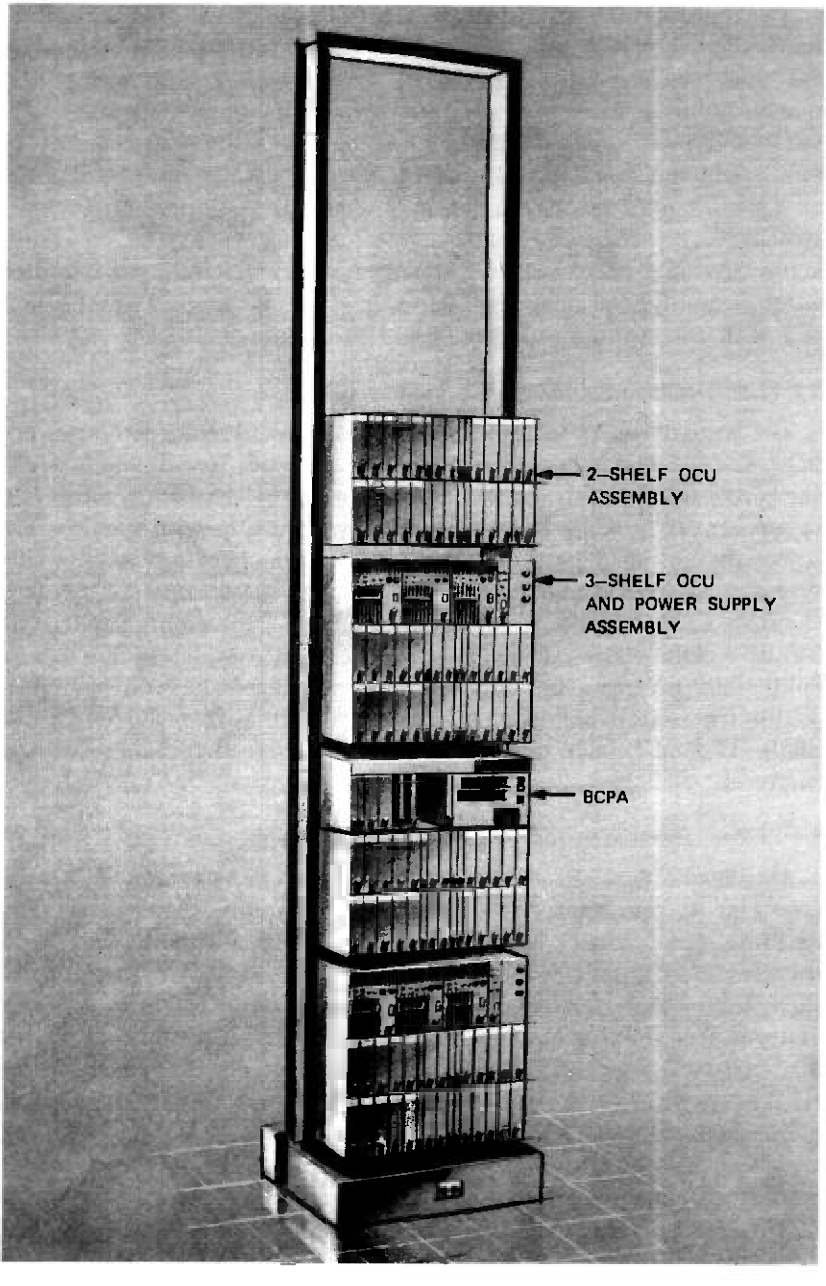

Fig. 3—Standard arrangement for office channel unit bay.

3.2 Subrate-data-multiplexer arrangements

The subrate-data-multiplexer (SRDM) subsystem comprises one BCPA shelf, one 5-V power supply shelf for distributing redundant 5-V power to the entire subsystem, one three-shelf SRDM and performance monitor assembly, and five two-shelf SRDM assemblies. The SRDM and performance monitor assembly, as well as each SRDM assembly, can accommodate 40 (at a 9.6-kb/s data rate) or 80 (at a 4.8- or 2.4-kb/s data rate) 64-kb/s, DS-0 channels. Accordingly, the subsystem has a maximum capacity of 480 DS-0 channels. The subsystem is provided in a 7-foot two-bay arrangement. By deleting one two-shelf SRDM assembly, the subsystem is provided in one single 11-foot 6-inch bay with a maximum capacity of 400 DS-0 channels.

3.3 T1 data-multiplexer, large-office arrangements

The large-office, T1 data-multiplexer (T1DM) subsystem comprises one BCPA shelf; one T1 data-multiplexer performance monitor (T1DM–PM) shelf; and four four-shelf T1DM assemblies, three of which contain four in-service T1DMs while the fourth contains three in-service T1DMs and one T1DM arranged as the spare for the other 15 T1DMs in the subsystem. Each of the in-service T1DMs can accommodate 23 DS-0 channels and, accordingly, the subsystem has a maximum capacity of 345 DS-0 channels. This subsystem is provided completely in a 7-foot, double-bay arrangement. By deleting one four-shelf T1DM assembly, containing four in-service T1DMs, the subsystem is provided in one single 11-foot 6-inch bay with a maximum capacity of 253 DS-0 channels.

3.4 T1 data-multiplexer, local-office, initial-bay arrangements

The local-office, initial-bay subsystem for offices employing T1DMs (see Fig. 4) comprises one four-shelf local timing supply and T1DM assembly (containing a local timing supply with a subset of the BCPA features, one in-service T1DM and one spare T1DM), one T1DM–PM shelf, three single T1DM shelves for in-service use, two three-shelf OCU and power supply assemblies, one two-shelf OCU assembly, four (one for each in-service T1DM) multiplexer jack-and-connector panels (M-JCPs), and three (one for each OCU assembly) subrate-data-multiplexer jack-and-connector panels (SM-JCPs). Each in-service T1DM can accommodate 23 DS-0 channels and each OCU assembly can accommodate 20 customer channels. Accordingly, the subsystem has a maximum capacity of 92 DS-0 channels and 60 customer channels. This subsystem is provided completely in a 7-foot two-bay arrangement. By deleting one three-shelf OCU and power supply assembly

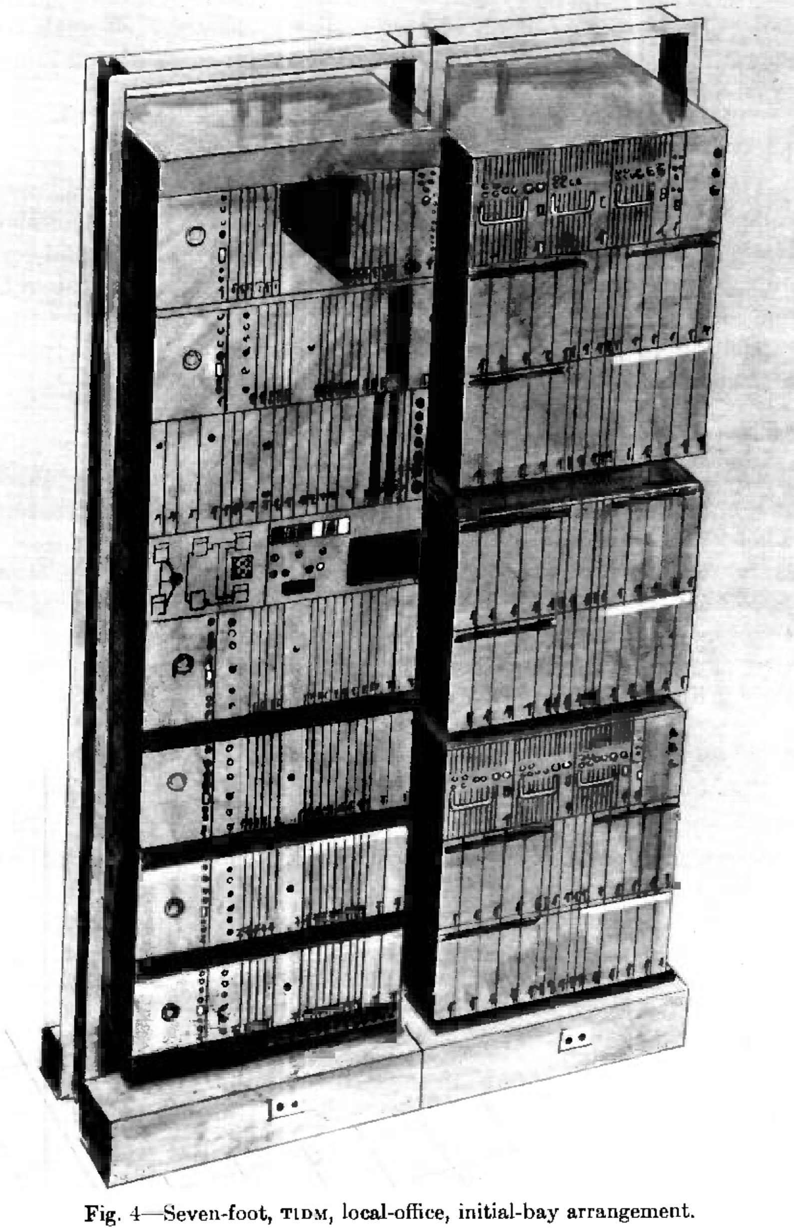

Fig. 4—Seven-foot, TIDM, local-office, initial-bay arrangement.

and its associated SM-JCP the subsystem is provided in one single 11-foot 6-inch bay with a maximum capacity of 92 DS-0 channels and 40 customer channels. As in the case of the OCU subsystem, when employed in local offices, 5- and 10-channel integral subrate multiplexing can be used so that if all of the available customer channels were operating at subrate speeds, only one-fifth to one-tenth of that number of DS-0 channels would be required.

3.5 T1WB4 data-voice-multiplexer hub-office arrangements

The T1WB4 data-voice-multiplexer (T1WB4) hub-office subsystem comprises one BCPA shelf and seven two-shelf T1WB4 assemblies. Each of the seven T1WB4s contains its own monitoring and spare multiplexing circuitry, and accommodates 12 port circuits providing one DS-0 channel each. Accordingly, the subsystem has a maximum capacity of 84 DS-0 channels and is provided in either a 7-foot, two-bay arrangement or in one single 11-foot 6-inch bay.

3.6 T1WB4 data-voice-multiplexer local-office arrangements

The local-office subsystem for offices employing T1WB4s comprises one BCPA shelf; one two-shelf T1WB4 assembly with an associated M-JCP; and either additional T1WB4 assemblies with one M-JCP for every two T1WB4s or a number of three-shelf OCU and power supply assemblies and two-shelf OCU assemblies, each with associated SM-JCPs. The exact quantities of additional T1WB4 assemblies or OCU assemblies depend on whether the subsystem is provided in 7-foot bays or 11-foot 6-inch bays and on exactly which option is chosen.

3.7 Nodal or secondary timing-supply arrangements

The nodal or secondary timing-supply (NTS or STS) subsystem comprises nothing more than one NTS or one STS in either a single 7-foot bay or a single 11-foot 6-inch bay. It is considered as a separate subsystem because, after obtaining reference timing signals from two DS-1 channels entering the office, its functions are completely independent from all the other subsystems in that office except that it supplies timing to them.

3.8 Multipoint-junction-unit arrangements

The 11-foot 6-inch bay arrangement for the multipoint-junction unit (MJU) comprises one BCPA shelf, and four two-shelf MJU assemblies with two associated power supply shelves. In the 7-foot, two-bay arrangement, six two-shelf MJU assemblies and three associated power supply shelves can be provided in addition to the BCPA shelf.

Each two-shelf MJU assembly may be equipped with up to sixteen full duplex (FDX) MJUs. An FDX MJU consists of either one or two identical circuit packs. When one circuit pack is utilized, a three-branch, multipoint circuit is provided. With the addition of the second circuit pack, a five-branch, multipoint circuit is formed.

3.9 Serving-test-center arrangements

The serving-test-center (STC) subsystem comprises a number of testboards for DS-0 channel test access and a number of cross-connect bays with complete flexibility for interconnecting the DS-0 channels from the various subsystems through the testboards.

3.9.1 Testboard

The testboard (Fig. 5) is the heart of the STC. It contains the test equipment needed to troubleshoot and maintain DDS circuits. A strong concern in the physical design of this unit was ease of operation by test center personnel.

There are two major sections to the testboard: a jack field for accessing individual DDS circuits and a control panel with writing shelf.

The jack field is composed of horizontal panels each providing 30 jack modules. Each module has a designation pin color-coded according to the particular customer's data rate. In many cases, by using this feature a craftsperson can troubleshoot the circuit without obtaining a circuit-record card. The jack module containing six miniature jacks represents a significant size reduction over existing central office jacks. Its small size permits up to 450 customers' circuits to appear in one testboard. Each jack module may be replaced individually.

The control panel has a sloping face that accommodates the portable digital test sets as well as other control units associated with generating test codes and establishing voice communications with the testboard. A writing shelf is provided for the test operator's use.

The emphasis in the physical design of the digital test sets was placed on arranging indicators, keys, switches, and designations in a logical and readily usable fashion. Figure 6 shows the test sets and the power and signal cords housed therein. Both sets use identical cases, framework and brackets, power supplies, and cords. The structure consists of two aluminum extrusions mounted on the sides that provide mechanical support for the printed boards, and a formed aluminum shell and cover. The sets may be completely disassembled and yet remain electrically operative for diagnosing trouble in their operation.

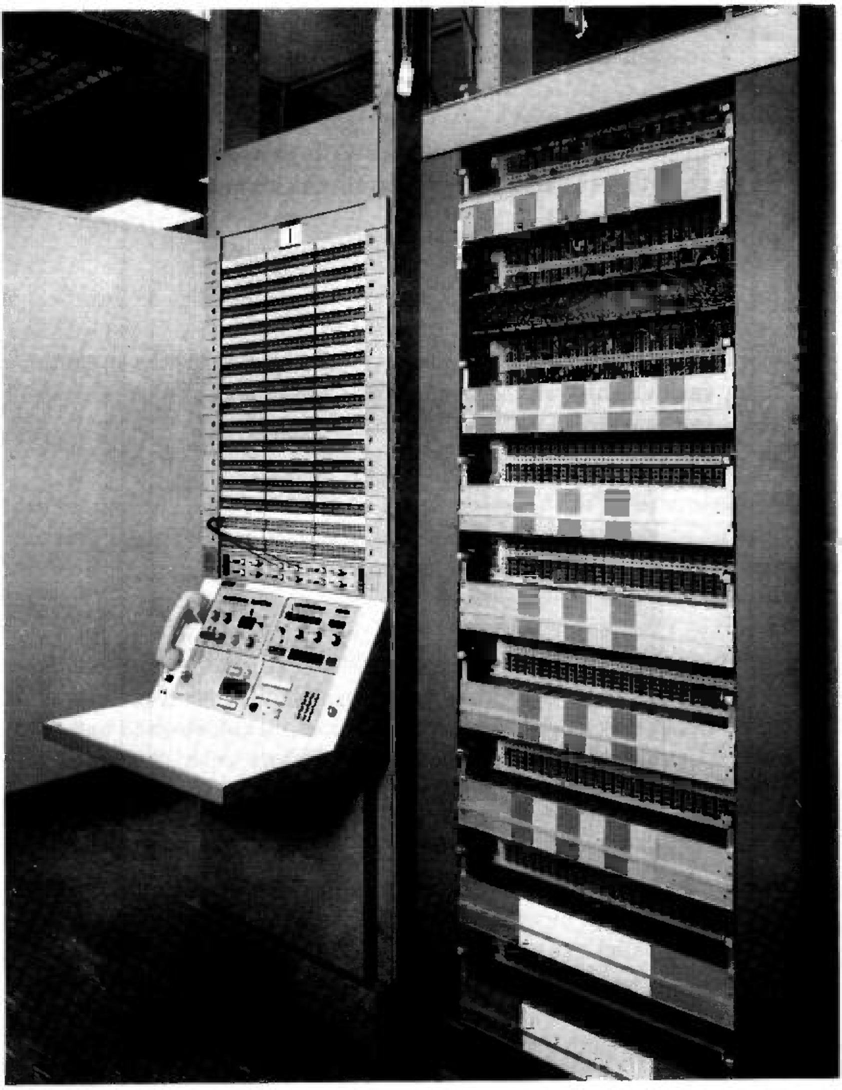

Fig. 5—950A testboard and cross-connect bay.

3.9.2 Cross-connect bay

To provide flexibility in interconnecting DDS equipment, all signals at the DS-0 level (64 kb/s) pass through a cross-connect field. This also provides access to each customer's circuit for the testboard. Constraints on the physical design included tight performance specifications on the plug-in contacts, ease of making and rearranging connec-

Fig. 6—Digital transmitter and receiver test sets.

tions, and provision for growth as the number of customers served by a particular office grows.

A versatile cross-connect system was developed that uses permanently wired, rear terminations from the various DDS equipment and quick-change, pluggable, front-panel connections for arranging the individual circuits. Each cross-connect panel (see Fig. 7) contains 400 quad (four-wire) terminations which accept latching quad jumpers. To eliminate patching errors and to assure optimum reliability, this patching system features polarized plugs with protected contacts for "hit-free" operation. The jumper plug housings have built-in latches that assure firm retention when installed in the panel, yet allow easy removal with a simple tool. The plugs are color-coded for fast, positive identification of jumper length and have snap-in, stamped-and-formed contacts that offer minimum cost and maximum serviceability.

The jumper contacts have a cantilever-beam engagement spring to insure controlled contact pressures with minimum wear on the gold-plated surfaces. An external retention spring provides quick assembly and firm seating in the jumper plugs, allowing on-site assembly of nonstandard length jumpers and the ability to repair defective units.

Each panel is equipped with a plastic wiring duct that functions as a fanning strip and provides a path for routing the quad jumpers. A hinged duct cover, which keeps the jumpers in place and has designation cards on both sides for equipment-termination records, is provided with each panel.

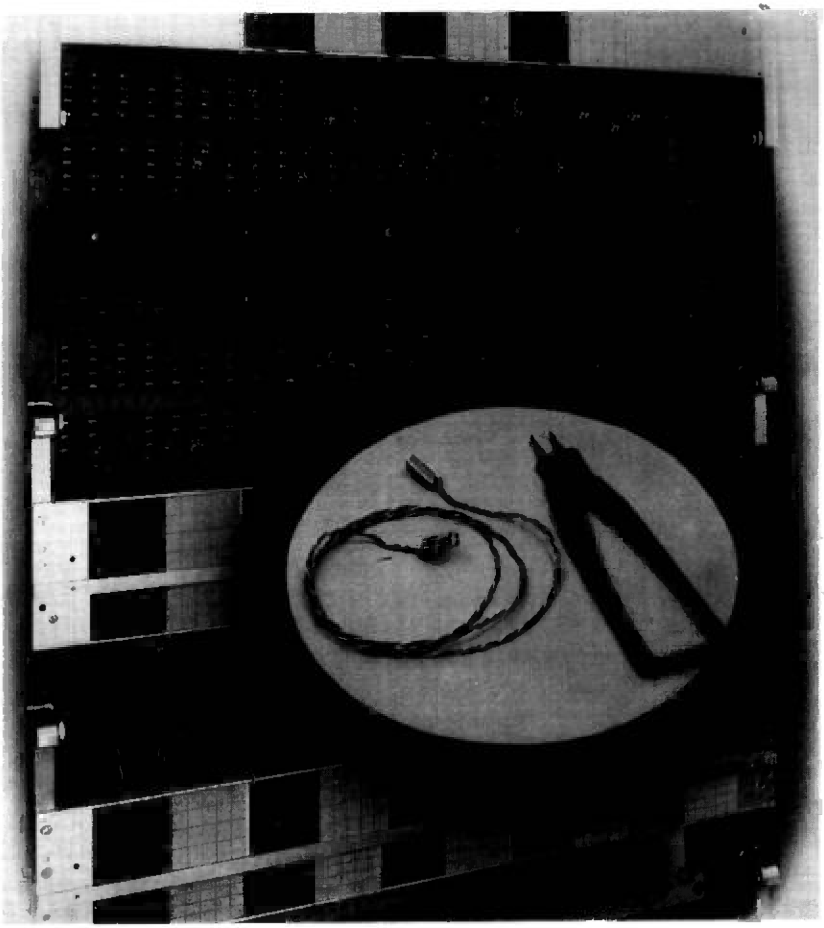

Fig. 7—Cross-connect panel, jumper, and removal tool.

The panels are combined in a cable-duct type bay (see Fig. 5). Filler plates that support wiring ducts increase the size of the bay, resulting in a typical width of approximately 42 inches.

The bay is equipped with one to nine quad-terminal-panel assemblies, according to job requirements, which provide a maximum of 3600 quad terminations. The quad terminals on each panel are arranged in two groups of 4 rows and 50 columns. Based on a random interconnection of equipment and a normal distribution of jumper wires in the horizontal ducts, the DSX-0 cross-connect is limited to six bays to ensure that the wire-handling capacity of the ducts is not exceeded.

3.10 Office arrangements

The several subsystems described above are utilized in various combinations to form the two basic types of DDS office configurations: the hub office and the local office. The local offices each "home" directly or indirectly on a given hub office. Each subsystem contains equipment that is physically independent from that contained in any other subsystem within the limitations imposed by certain cable-length requirements. Locations of the various 11-foot 6-inch bays or the 7-foot two-bay combinations in any given office may vary because of restrictions imposed by the location of related equipment such as the main distribution frame, the DSX-1 cross-connect, or, in hub offices, the DSX-0 cross-connect and its associated testboards.

3.10.1 Hub offices

A hub office is an office in the DDS that combines the T1 data streams from a number of local offices into signals suitable for transmission over DDS facilities, and provides test access by means of an STC subsystem. Hub offices, in general, will have equipment functioning in four different sections.

(*i*) Serving-test-center cross-connect section.
(*ii*) Long-haul-access multiplexing section.
(*iii*) Local-access multiplexing section.
(*iv*) End-access section.

The STC cross-connect section comprises testboards, multipoint-junction-unit subsystems and at least the STC cross-connect portion (DSX-0A) of the DSX-0. The testboards and DSX-0A function together and are in close physical proximity. The multiplexer cross-connect (DSX-0B) functions both with the long-haul-access equipment and with the local-access equipment. Accordingly, the DSX-0B may be split and each portion placed with its associated equipment, or, if a centralized cross-connect area is more desirable and cabling restrictions permit, the DSX-0B may be placed contiguous to the DSX-0A.

The long-haul-access multiplexing section comprises T1DM subsystems, SRDM subsystems, and the nodal or secondary timing supply subsystem. As mentioned above, the required DSX-0B cross-connect bays may be located within this section. The local-access multiplexing section comprises T1DM subsystems, T1WB4 subsystems, and SRDM subsystems. Again, the required DSX-0B cross-connect bays may be located in this section. The end-access section comprises only OCU subsystems.

3.10.2 Local offices

A local office is an office in the DDS that passes to a hub office circuits that enter the building over local loops or over T1 lines from other local offices. Such an office will usually have equipment functioning in only one section. It is expected that local offices will be small initially. Accordingly, the subsystems that make up the initial bay or bays comprise combinations of various equipments to make efficient use of bay space. A local office employing T1DM service may initially have a T1DM local-office, initial-bay subsystem. A local office employing T1WB4 service must have a T1WB4 local-office subsystem.

Additional OCU growth is obtained by the installation of OCU subsystems. When more efficient multiplexing than that available with the 5- or 10-channel integral subrate multiplexing accommodated in OCU assemblies is required, SRDM subsystems are employed. The DSX-0 cross-connect and test-access functions may be performed by multiplexer and subrate-data-multiplexer jack-and-connector panels, rather than by cross-connect bays and testboards as in hub offices.

IV. STATION APPARATUS

Two types of station apparatus are available to interface with the customer for transmission and reception of digital data over DDS facilities. If the customer requires an EIA interface for 2.4-, 4.8-, or

Fig. 8—Data service unit.

Fig. 9—Channel service unit.

9.6-kb/s operation or a CCITT interface for 56-kb/s operation, a DSU (see Fig. 8) is provided. If the customer requires only the bipolar signal as received from the local loop, a CSU (see Fig. 9) is provided.

Since these units are located on the customer's premises, they play a major role in determining how the customer views DDS. They must be easy to install and maintain, perform well, and present a favorable appearance. In terms of physical design, factors such as styling, size, and human engineering are particularly important. Both the DSU and CSU relate closely in these factors to the new family of analog data sets being offered by the Bell System. All are packaged in low-profile aluminum or aluminum-finished housings with black molded covers, and provide status indicators to help the customer monitor system operation.

4.1 Data service unit

The circuitry for the DSU is partitioned into the following functional blocks: customer interface, transmit logic, line driver, line receiver, timing recovery, receive logic, test circuits, and power supply. To minimize the number of circuit pack codes and interconnecting wires, two basic circuit boards were designed. The logic board, 70 square inches in area, contains the transmit and receive logic. It is used in all DSUs regardless of data rate. The analog board, 100 square inches in area, contains, with the exception of the power supply, the balance of the circuitry, including the interface connector, a test switch, and light-emitting-diode indicators (LEDs). Since the analog board is speed dependent, four coded circuit packs are manufactured from the same basic board design. The power supply, utilizing a ferroresonant transformer, is packaged separately in a metal enclosure to eliminate high-voltage exposure on the two circuit boards.

For the overall unit, a low-profile shape was chosen that allows stable stacking of several DSUs. The dimensions of the unit are approximately $11\frac{1}{2}$ inches wide, 4 inches high, and $10\frac{1}{2}$ inches deep. It weighs approximately 10 pounds.

As shown in Fig. 10, the circuit boards are arranged horizontally. Service option switches are located at the edges of the boards for installer access. Special design consideration was required in arranging the boards and power supply and in selecting and placing components to ensure proper operation in a maximum 120°F operating ambient.

Tooling costs and assembly operations were reduced by the design of two aluminum extrusions that provide the circuit-board guides,

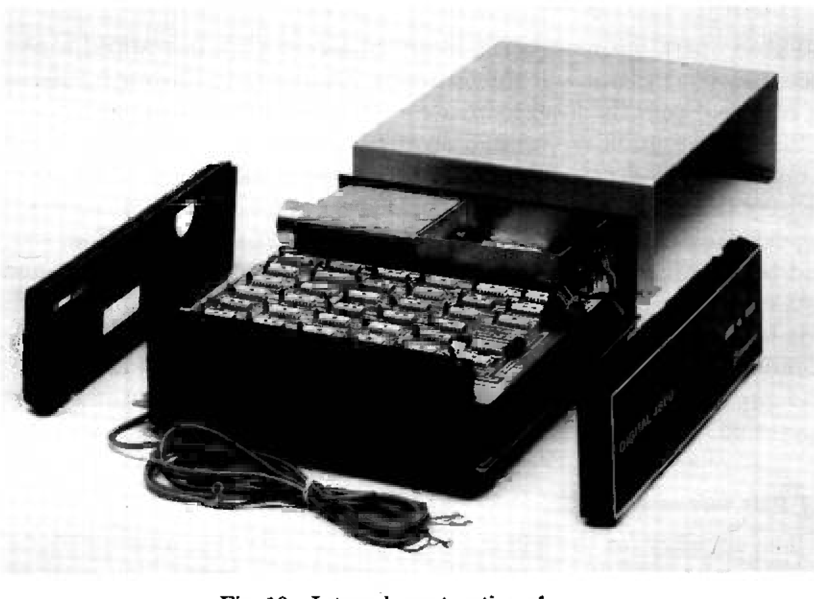

Fig. 10—Internal construction of DSU.

power supply shelf, and mounting rails, as well as the basic chassis assembly. The front and rear covers, injection molded of clear polycarbonate and back-painted black, snap onto the chassis and brushed aluminum housing. Both covers accept the test switch and indicator panel, which is connected by cable to the analog board. This allows the customer controls to be placed on the same or opposite surface as the interface connector and power cord.

DSUs may be stacked and mechanically fastened together in arrangements of up to three sets. When larger numbers are required, a cabinet arrangement for up to ten DSUs is provided. In this case, the units are mounted vertically on edge with the aluminum housing removed for maximum heat dissipation.

4.2 Channel service unit

For the CSU, study of circuit partitioning yielded four basic functional blocks: transmitter, receiver, test circuit, and power circuit. Since no customer control is required, a wall-mounted package which occupies no table or cabinet space is felt to be the most desirable method of packaging for customer convenience, although care must be taken to occupy the least amount of wall space. The basic CSU structure is a nest of two circuit boards, each approximately 35 square inches in area, assembled with components facing each other. One circuit board contains the receiver circuitry and the other the transmit, power, and test circuitry. A fully shielded shunt-type ferroresonant transformer provides the necessary DC voltages on either board. Both boards are speed dependent, so a total of eight circuit pack codes are required.

The housing is a two-piece injection-molded polycarbonate shell. The top is molded of clear material and back-painted black; the base is finished with aluminum paint to give the same material appearance as other new Bell System data station sets. The entire package measures approximately 8 inches wide, 5 inches high, and $2\frac{3}{4}$ inches off the wall.

Multiple arrangements of up to 20 CSUs can be provided in a cabinet.

V. CONCLUSION

The physical design of DDS equipment is a response to the performance and economic challenges of this new system. Central office hardware emphasizes flexibility and provision for growth, while station sets are marked by modern styling and ease of operation. Future designs will build upon these bases as new digital data communications services are offered.

VI. ACKNOWLEDGMENTS

The work outlined here was a joint effort, combining the talents of every individual in the Data Systems Physical Design Department. Each task, whether in the area of design, technical investigation, documentation, manufacturing, or field support, was an important constituent. Several who contributed to the job have moved to other areas and two individuals who played important roles, Norman T. Rauch and John P. Slickers, are deceased. Their contribution and the labors of all who participated in the physical design effort are gratefully acknowledged.

Contributors to This Issue

E. Craig Bender, B.S.E.E., 1964, Syracuse University; M.S.E.E., 1966, University of California at Los Angeles; Bell Laboratories, 1966–1974; AT&T, 1974—. At Bell Laboratories, Mr. Bender was primarily involved in system studies of wideband data services. He was most recently concerned with the engineering and performance aspects of local DDS transmission and is currently an Assistant Engineering Manager with AT&T working on new data network services.

Paul Benowitz, B.S., 1952, Brooklyn College; B.S.E.E., 1958, Polytechnic Institute of Brooklyn; M.S.E.E., 1961, New York University; Bell Laboratories, 1957—. Mr. Benowitz has worked on the conversion of TWX service from a private line network to an integral part of the DDD network. He also participated in the development of the 43B1 Voice Frequency Carrier Data System, the 1A Data Station, and the Data Line Concentrator System. He supervised the development of multiplexers for the Digital Data System and is currently supervising investigation of message switching applications.

P. F. Brown, Jr., B.S.E.E., 1965, Duke University; M.S.O.R., 1966, Cornell University; 1971, Columbia University; Bell Laboratories, 1965–72; Southern Bell Telephone and Telegraph Company, 1972—. Mr. Brown initially worked in data station engineering and studies of communications for business information systems. He later was involved in studies of buffered digital networks and planning for the Digital Data System.

S. J. Butterfield, B.S.E.E., 1967, Ohio University; S.M.(E.E.), 1968, Massachusetts Institute of Technology; Bell Laboratories, 1967–1973; New England Telephone, 1973—. At Bell Laboratories, Mr. Butterfield was engaged in system studies and equipment design for the L-mastergroup digital system, T1, T1C and coaxial digital systems, and the T1WB4 data-voice multiplexer for the DDS. Mr. Butterfield is currently concerned with T1 Carrier performance with the New England Telephone Company. Member, IEEE, Phi Kappa Phi, Sigma Xi, Eta Kappa Nu, Tau Beta Pi.

Michael P. Cichetti, Jr., B.E.E., 1966, Cooper Union, New York; M.S.E.E., 1967, Stanford University; Bell Laboratories 1966—. Mr. Cichetti has worked on low-speed data transmission and low-speed data multiplexing prior to his work on DDS multiplexers. He is currently involved with interactive message switching. Member, Tau Beta Pi, Eta Kappa Nu.

Thomas G. Cross, B.S.E.E., 1963, California State Polytechnic College; M.S.E.E., 1965, Northwestern University; Bell Laboratories, 1963–1968, 1970–1972; American Telephone and Telegraph Company, 1968–1970, 1973—. At Bell Laboratories, Mr. Cross was involved in development of long-haul microwave radio systems, short-haul digital carrier systems, and wideband digital data terminals. At AT&T, he was involved in long-range fundamental planning for transmission systems and is currently responsible for technical regulatory matters.

S. M. Fitch, B.S.E.E., 1960, and M.S.E.E., 1962, Purdue University; Ph.D., 1973, New York University; Bell Laboratories, 1963—. Mr. Fitch initially worked on the development of data stations for the No. 1 ESS ADF Data Switching System and then on the 85 and 86 Data Selective Calling Systems. More recently, his work involved developing multipoint services for the DDS. He is presently supervisor of the Data Network Testing Group and is responsible for developing maintenance and testing equipment for DDS. Member, IEEE, Eta Kappa Nu.

H. C. Illium, B.S.A.E., 1954, Purdue University; M.E.E., 1960, New York University; Bell Laboratories, 1956—. Mr. Illium has been involved with the physical design of data sets and data transmission equipment. He currently supervises a data systems physical design group at Holmdel. Member, Sigma Gamma Tau.

Nicholas Knapp, Jr., B.S.E.E., College of the City of New York, 1929; Bell Laboratories, 1929–1973. Mr. Knapp was engaged largely in physical developments. His early work was on toll switching systems, switchboards, toll ticket distributing systems, and radio. During World War II he worked on land-based, shipboard, and airborne radar equipment. Subsequently, he was engaged in coaxial and military carrier systems and on data sets and data systems.

Joseph G. Kneuer, B.S.E.E., 1958, University of Notre Dame; M.S.E.E., 1965, and Ph.D., 1966, Syracuse University; Bell Laboratories, 1966—. At Bell Laboratories, Mr. Kneuer has worked in high-speed data communication and later in the development of the Digital Data System. Since 1972 he has supervised a circuit design group responsible for the local distribution system of DDS. Member, IEEE, Sigma Xi.

William J. Lawless, B.S.E.E., 1959, Newark College of Engineering; M.S.E.E., 1961, New York University; Ph.D., 1971, Polytechnic Institute of Brooklyn; Bell Laboratories, 1959—. Mr. Lawless has worked on the design and development of data communication equipment. He is currently head of the Data Network Transmission Department. Member, IEEE, Tau Beta Pi, Eta Kappa Nu.

W. B. Lueft, B.S.E.E., Worcester Polytechnic Institute, 1953; Western Electric, 1953–1961; Bell Laboratories, 1961—. Currently, Mr. Lueft is supervisor of a physical design data testing group. Senior member, IEEE; member, Eta Kappa Nu.

J. J. Mahoney, Jr., A.T.&.T., 1926–34; Bell Laboratories, 1934—. At Bell Laboratories, Mr. Mahoney has worked on studies of electrical protection of telephone facilities, the design of oscilloscopes used in development of radar systems, and on systems engineering studies for carrier and wideband data transmission. He worked on performance objectives for the Digital Data System and currently supervises a group responsible for data performance studies. Senior Member, IEEE.

James J. Mansell, B.S.E.E., 1962, St. Louis University; M.S.E.E., 1963, University of Pennsylvania; Bell Laboratories, 1962—. Mr. Mansell has been involved in systems engineering of data communications networks. He currently supervises a data systems operations group.

Richard C. Matlack, B.S., 1938, University of Pittsburgh; U. S. Army, 1941–1945; Bell Laboratories, 1945—. Mr. Matlack has been involved in switching development of the No. 5 Crossbar systems including the planning for direct dialing arrangements for the step-by-step system and concentrator arrangements. While in the Army, he was engaged in radar and electronic gun director work with the Army Antiaircraft Artillery Corps. He is past secretary-treasurer of the Professional Group of Electronic Computers; member, AIEE, Eta Kappa Nu, and Sigma Tau.

G. W. Phipps, B.S.E.E., 1960, Rose Polytechnic Institute; M.S.E.E., 1963, New York University; Bell Laboratories, 1961—. Mr. Phipps initially worked in data station engineering for both line-switched and store-and-forward data systems. He supervised a group responsible for studies of communications for business information systems and later a group responsible for computer-access planning for *Picturephone*® service. He was involved in network planning for the Digital Data System and software to aid telephone companies in planning and administering the network. Mr. Phipps is currently head of a department responsible for data message switching planning and software development. Member, IEEE, Tau Beta Pi.

D. L. Rechtenbaugh, B.S.E.E., 1960, South Dakota State College; M.E.E., 1962, New York University; American Telephone and Telegraph Company, 1969–1972; Bell Laboratories, 1960–1969 and 1972—. Mr. Rechtenbaugh initially worked on maintenance planning for private line and *DataPhone*® services and the No. 1 ESS ADF Data Switching System. During his assignment at AT&T, he was responsible for trunk maintenance planning including centralized automatic testing arrangements. He is presently supervisor of the Digital Network Maintenance Planning Group and is responsible for developing the T-Carrier Administration System and maintenance planning for the emerging digital network.

Donald W. Rice, B.S.E.E., 1958, University of Pennsylvania; M.S.E.E., 1959, California Institute of Technology; Bell Laboratories, 1959—. Mr. Rice has been engaged in various development activities on data systems including systems planning, circuit design, and physical design. Member, IEEE, Eta Kappa Nu, Tau Beta Pi.

Burton R. Saltzberg, B.S.E.E., 1954, New York University; M.S., 1955, University of Wisconsin; Eng. Sc.D., 1964, New York University; Bell Laboratories, 1957—. Mr. Saltzberg has been engaged in the development, exploration, and analysis of data transmission techniques. He presently supervises a group involved in the development of equipment for data communication networks. Member, IEEE, Eta Kappa Nu, Tau Beta Pi, Sigma Xi.

Norman E. Snow, B.S.E.E., 1952, University of Arkansas; Southwestern Bell, 1952–1957; Bell Laboratories, 1957—. Mr. Snow initially worked in the Arkansas area of Southwestern Bell where he engineered toll terminal, local switching, and central office power installations. At Bell Laboratories, he has worked in both systems engineering and development of voiceband and wideband data sets and systems for a variety of standard and special-purpose Bell System data services. Registered professional engineer; member, Tau Beta Pi, Theta Tau, Phi Eta Sigma, Pi Mu Epsilon.

L. A. Spindel, B.S.E.E., 1967, Lehigh University; M.S.E.E., 1969, Rutgers University; Bell Laboratories, 1967—. Mr. Spindel has worked on the systems engineering aspects of various data communications services, including analog data transmission and network planning for the Digital Data System. He is presently supervisor of an analytical support group for marketing and economic analyses of new data services. Member, IEEE, Phi Beta Kappa, Tau Beta Pi, Eta Kappa Nu, Omicron Delta Kappa.

Herbert M. Zydney, B.A., 1954, Columbia College; B.S.E.E., 1955, and M.S.E.E., 1959, Columbia University; Bell Laboratories, 1959—. Mr. Zydney has been engaged in data communications systems and the design of circuits and computer programs for digital communications in the Bell System. He was responsible for a number of developments for the Digital Data System. He is presently Director of the Facilities Network Planning Center, which provides planning methods for the evolution of the facilities used in telephone service and improved procedures for the activities needed in operating the telephone network. He is a registered professional engineer. Member, IEEE, Tau Beta Pi, Eta Kappa Nu.