THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 60, No. 6, July–August 1981
Printed in U.S.A.

No. 4 ESS:

Digital Interface

By K. M. HOPPNER, H. MANN, S. F. PANYKO, and J. VAN ZWEDEN

(Manuscript received August 22, 1980)

This article describes the hardware and maintenance software implementation used for the Digital Interface (DIF), a new transmission interface for the No. 4 ESS. The DIF replaces the Digroup Terminal (DT) and the Signal Processor 2 (SP2) used for terminating digital carrier trunks. It also provides a more economical transmission interface for the No. 4 ESS, enhances the use of the LT-1 connector for terminating analog carrier facilities, and provides a standard peripheral/processor interface for maintenance. In conjunction with the DIF, a hierarchical modularly structured maintenance software system was introduced to support the new No. 4 ESS peripherals. Incorporated into this system were the mechanisms required to support the introduction of microcomputer-based peripherals such as DIF.

I. INTRODUCTION

The Digital Interface (DIF) is a newly introduced No. 4 ESS peripheral frame whose functions combine the operations performed by the Digroup Terminal (DT) and the Signal Processor 2 (SP2), with the exclusion of the supplementary matrix frame which is an SP2 optional adjunct.[1,2]

The main functions of a DIF are to terminate DS-1 level signals and to multiplex them to a form suitable for the No. 4 ESS digital switch; to perform the necessary signaling interchange between the transmission and switching facilities; and to provide adequate fault detection and reconfiguration capability. These functions are identical to those performed by the DT/SP2 complex but the DIF achieves them in a more compact and modern fashion. The DIF occupies less than one-half the

space and uses two-thirds less power than the DT. This reduction was achieved in a number of ways:

(*i*) by combining the equivalent of four DTs and an SP into a single structure thereby eliminating a number of duplicated functions and a complex interface,

(*ii*) by using custom and catalog large-scale integrated circuits instead of the 1A Technology previously used, and

(*iii*) by utilizing both metal oxide semiconductors and bit-sliced microcontrollers to replace the hardwired logic of the DT and SP controllers.

As will be discussed in more detail later, these changes resulted in a frame with an entirely different internal philosophy which was more flexible and more consistent with the architecture of the other peripheral frames in the No. 4 ESS. Sections II and III give a high-level logical and physical view of the DIF. The article then describes the Digital Interface Unit (DIU) architecture in Section IV, followed by the Digital Interface Controller architecture in Section V. Finally, the maintenance software developed for the DIF is presented in Section VI.

II. OVERALL FRAME ARCHITECTURE

A fully equipped DIF consists of a duplex controller (DIC), duplex Interface to the Peripheral Unit Bus (IPUB), 32 working DIUs, and two protection spare DIUs as shown in Fig. 1. The DIF connects directly to the Peripheral Unit Bus (PUB), interfaces directly or via an echo suppressor terminal with the Time Slot Interchange (TSI) of the No. 4 ESS network in the DS-120 format, and provides a DS-1 transmission interface with T1 facilities or LT-1 as shown in Fig. 2. Each DIU terminates five DS-1 signals, giving the DIF a capacity to terminate a

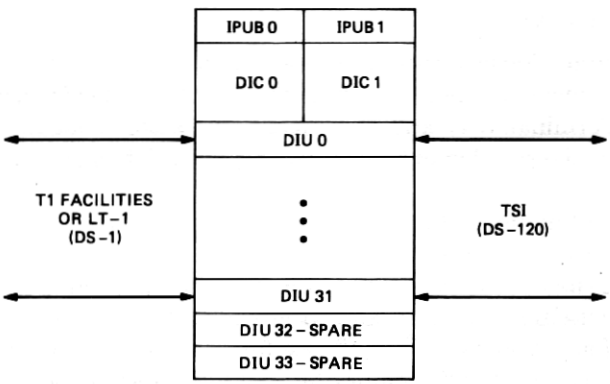

Fig. 1—DIF block diagram.

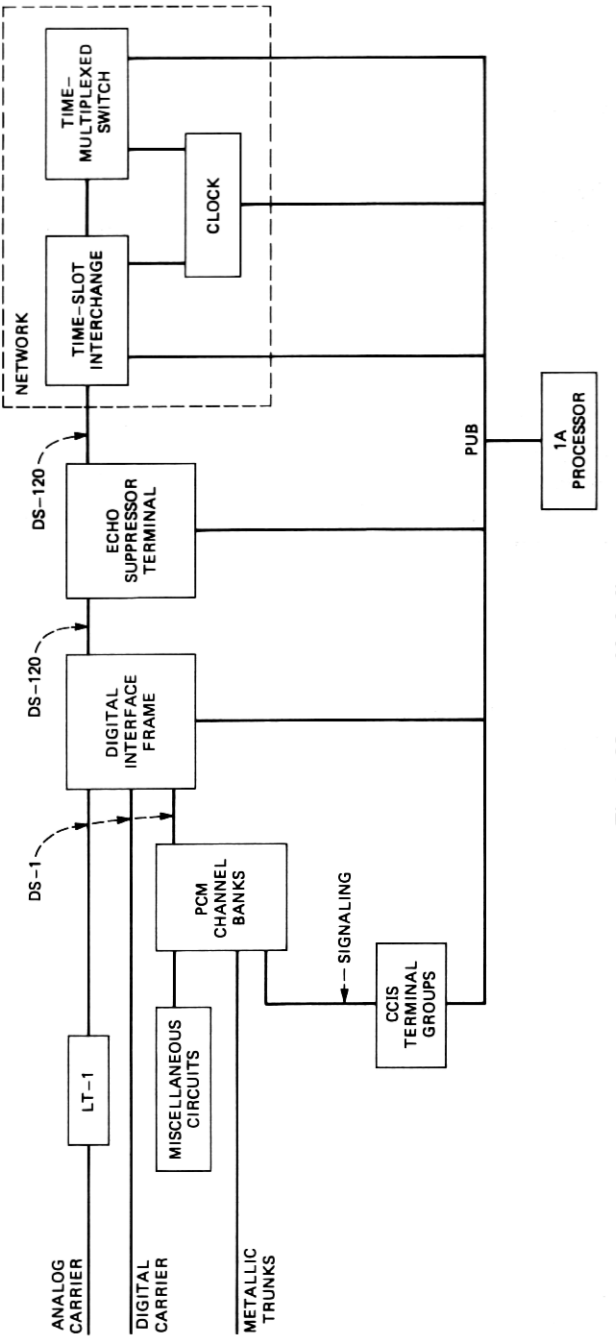

Fig. 2—No. 4 ESS block diagram.

total of 160 DS-1 signals (3840 trunks). A photograph of the DIF appears in Fig. 3.

III. PHYSICAL DESIGN

3.1 Overview

The circuitry which makes up the DIF is of two types. One type is of a transmission nature which is characterized by relatively few functional blocks arranged in a serial fashion. The other type is of a processor nature consisting of many functional blocks which have a high degree of interconnectivity typically in the form of large high-speed parallel buses. Within the limits of technology used in the DIF, circuit pack partitioning of the processor function indicated the need for a large Input/Output (I/O) circuit pack capability, whereas the transmission type of circuits required significantly less I/O. This difference influenced the physical design of the DIU, which is basically a transmission function, and that of the DIC, which is basically a processor or computer type of function.

The physical design objectives of the DIF included the following:

(*i*) Compatibility with the No. 4 ESS environmental, reliability, and frame I/O connector requirements.

(*ii*) Significant cost, space, and power reduction over the DT/SP2 complex which it replaces.

(*iii*) Circuit partitioning which facilitates interconnection, fault detection, fault diagnostics, and maintenance.

(*iv*) Physical embodiment which enables proper electrical functioning.

3.2 Device technology

The DIF uses low- and medium-power transistor-transistor logic, low- and high-power Schottky TTL, emitter-coupled logic, complementary metal-oxide semiconductor, and N-channel metal-oxide semiconductor integrated circuit technologies. The level or scale of integration ranges from small-scale integration to large-scale integration (LSI) with a fully equipped DIF using a total of about 18,000 devices of about 200 codes.

3.3 Frame and circuit-pack partitioning

The circuit-pack size chosen for the DIUs was determined by the DS-1 interface, described in Section IV, since this is the dominant entry in each DIU (five out of a total of nine circuit packs). Thus, a printed wiring board (PWB) which is nominally 8 inches high and 9 inches deep with a 114-lead connector was chosen for the DIU circuitry. This size made possible the double-sided PWB implementation (see Fig. 4) of the

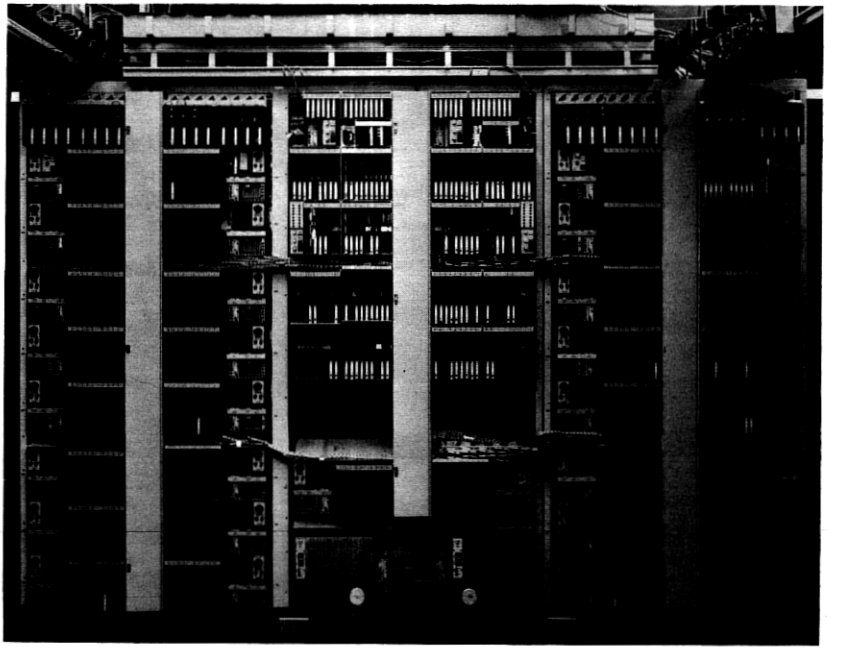

Fig. 3—Photo of DIF.

DS-1 interface which is the most cost-sensitive circuit in the DIF because of the large number used (170/DIF). The remaining DIU circuitry required two double-sided and two four-layer PWB circuit packs. The four-layer configuration is capable of accommodating approximately 50 percent more circuitry than the double-sided boards.

The same basic circuit-pack size is also used for the DIC; however, because of its large I/O requirements, a 184-lead connector is used instead of the 114-lead connector. The DIC consists of 94 circuit packs of which 48 are four-layer PWBs and the remaining 46 are double-sided.

Since all DIUs typically communicate with the controller via a bus, locating them about a central cable-duct (see Fig. 5) permits the bay cabling to be kept to a minimum length with a common access point for a pair of DIUs which minimizes the number of cable connectors required. Those circuit packs that interface with the DIC are located near the cable duct. Therefore, the circuit pack positions of the DIU on the right side of the duct are a mirror image of those of the DIU on the left side.

The DIF uses the No. 1 ESS framework: one double-bay frame which is 7 feet high, 6 feet 6 inches wide, and 12 inches deep; and one single-bay frame which is 3 feet, 3 inches wide. Two DIUs and their associated power units and switches are mounted on an 8-inch-high shelf. At the

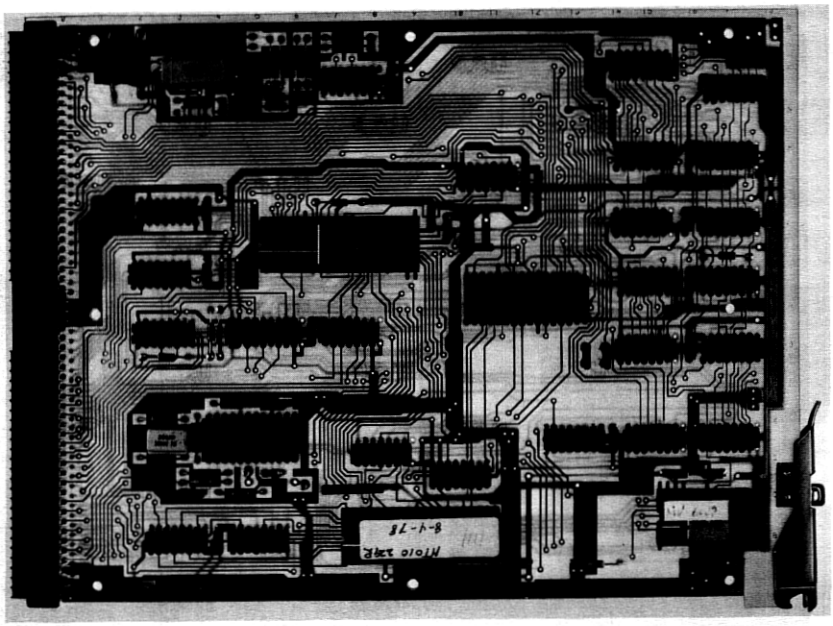

Fig. 4—Double-sided PWB.

top of bay 0 and bay 2 are the protection switch circuit packs and the DS-1 and DS-120 I/O connectors. These circuit packs are nominally 6 inches high and 9 inches deep. They contain the relays required for transferring the DS-1 and DS-120 signals from a DIU to a spare DIU. They also contain equalizers for the DS-1 lines and the transmit and receive coax connectors for the DS-120 signal. DS-1 connectors are located on a panel directly above the protection-switch circuit packs. The spare DIU in bay 0 provides backup protection for 15 DIUs located in bay 0 and the left DIU located in bay 1. Likewise, the spare DIU in bay 2 provides backup protection for the 15 DIUs located in bay 2 and the right DIU located in bay 1.

The duplex PUB interface is located at the top of bay 1 with its associated power units as shown in Fig. 5. The circuit packs which make up this function are nominally 4 inches high and 9 inches deep and use a 92-lead connector. The duplex controller is immediately below it and the associated power units are located at the bottom of the bay.

The circuit-pack positions of the PUB interface are not mirror-imaged about the cable duct as are those of the DIU and the DIC. This is for two reasons. First, the PUB interface does not use the cable duct for interconnecting to the controller. Second, the PUB is connected via cable connectors to the same backplane pins which the circuit packs

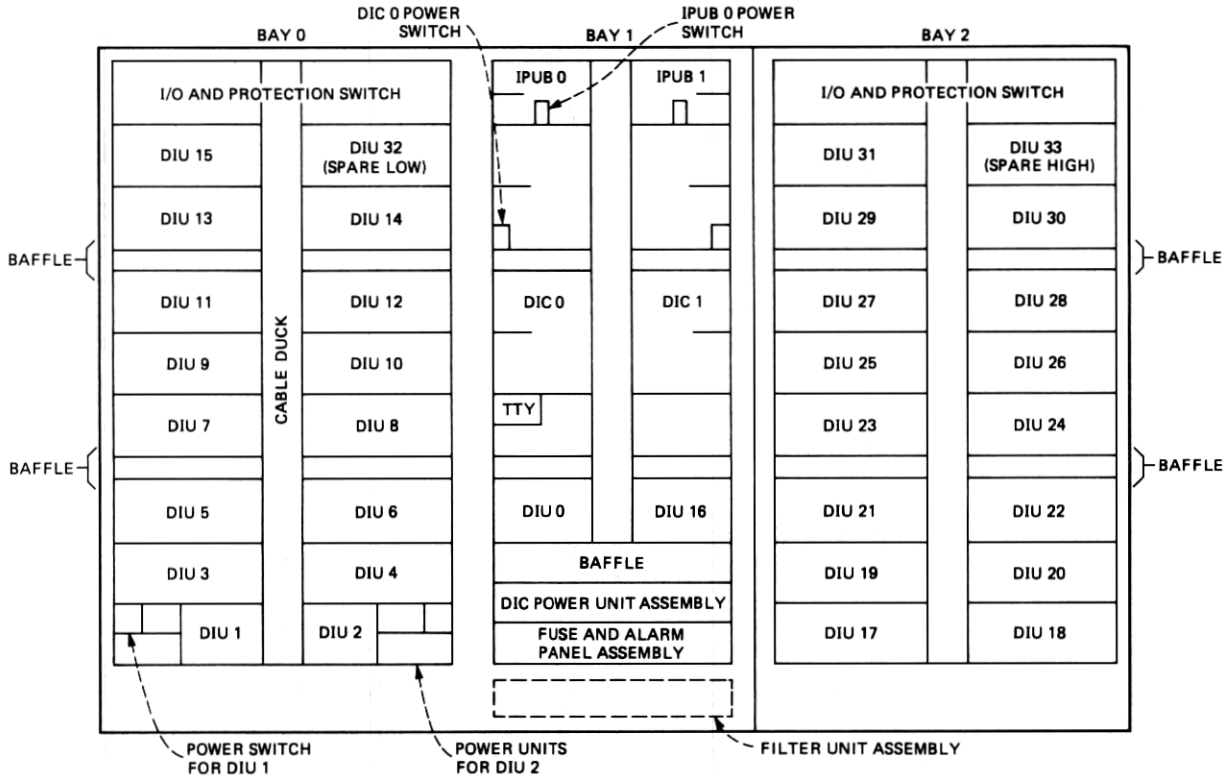

Fig. 5—DIF diagram (PUB interface and power units).

connect to; thus, to provide a PUB connector arrangement which is identical for both PUB0 and PUB1 to simplify installation, the PUB interface 0 and 1 circuit packs must have the same order. This arrangement is also used for other No. 4 ESS peripherals.

The power switches for the PUB interface, the DIC, and the DIUs are collocated with their respective functions to make their association self-evident.

Heat baffles are used to divert air flow out of the back of the frame and to cause aisle air to be taken in from the front of the frame at various levels. This prevents the cooling air from becoming excessively heated as it rises through the frame.

3.4 Frame cabling

Most of the frame cabling is 26 American wire gauge (AWG) twisted-pair, 24-conductor flat cable. The wires are on 0.062-inch centers with an untwisted section every 18 inches for terminating connectors (see Fig. 6). The cable is terminated to connectors which have insulation displacing terminals. These connectors plug onto 25-mil square pins located in the cable duct area.

The frame cabling is located on the circuit-pack side of the backplane

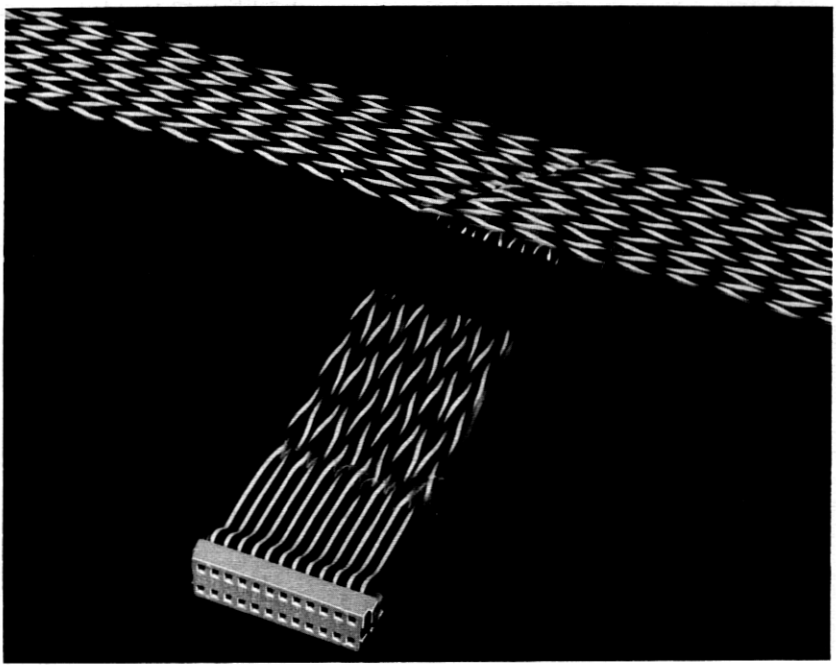

Fig. 6—Frame cabling (flat cabling.).

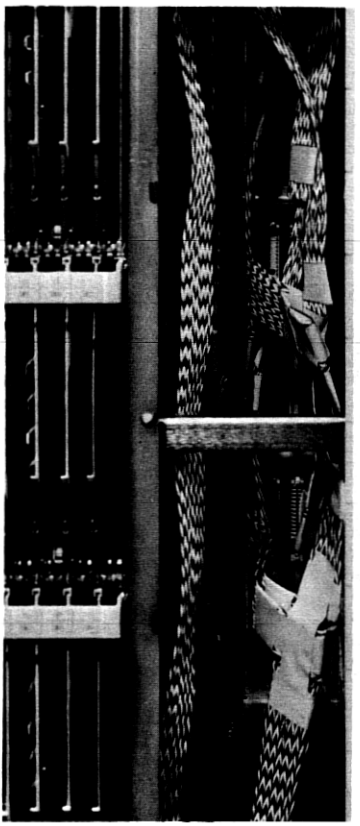

Fig. 7—Frame cabling (cable duct and local cable).

(see Fig. 7), and wiring to the cable connector pins is done in the backplane with surface wiring. This cabling scheme isolates the frame cabling from the backplane surface wiring by designating a specific area for frame cabling. The cable is routed horizontally between cable ducts in designated areas located in front of the frame upright on the top surface of heat baffles.

3.5 Circuit-pack design

The circuit-pack PWB sizes used in the DIF conform with the *Bellpac** standard circuit-pack size. The circuit-pack connectors and the latch used on these circuit packs are part of the *Bellpac* hardware system.[3] Before the design of circuit packs, design standards were determined to assure a functional and manufacturable design. These

* Trademark of Western Electric.

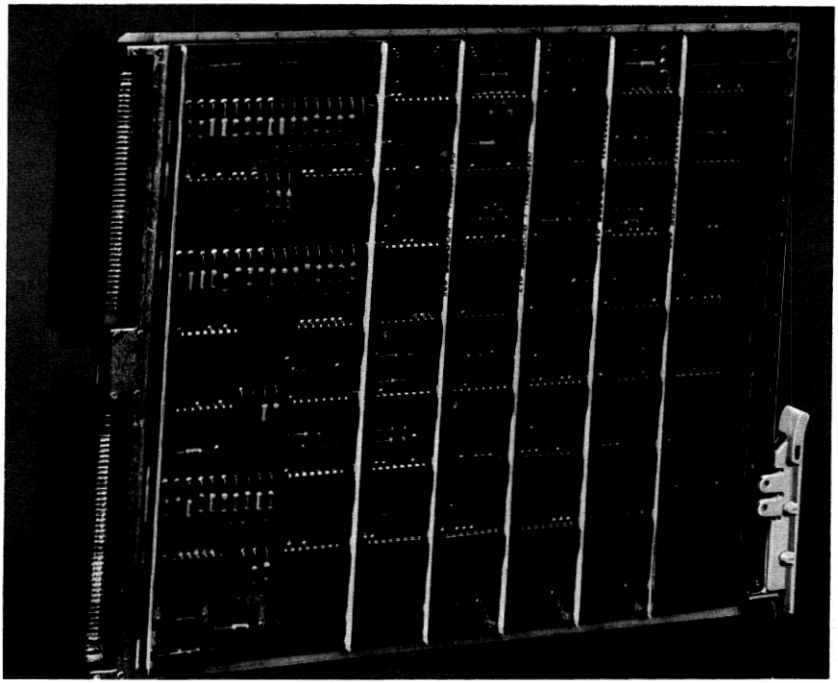

Fig. 8—DIC circuit pack utilizing appliqued bus bars.

standards include maximum component heights; component placement constraints; conductor path grid, sizes, and routing; plated-through hole sizes, lands, and placement; connector fanout pattern; placement of I/O devices; and use of power/ground decoupling capacitors.

An appliquéd bus bar is used on some double-sided PWB circuit packs for distributing power and ground (see Fig. 8). This makes more area available for routing signal paths, since wide printed paths are not needed for distributing power to the integrated circuits.

The four-layer PWB circuit packs consist of a power surface layer, two buried signal layers, and a ground surface layer (see Fig. 9). The signal layers are buried so that their fine features (metalization and clearances as small as 12 mils) are protected, whereas the power/ground layers which consist of relatively conservative features (metalization and clearances of 25 mils or greater) are placed on the outer layers.

3.6 Backplanes

The backplanes which interconnect the circuit packs consist of pin-populated PWBs—see Fig. 10. These backplanes use the *Bellpac* com-

pliant pin and are designed to *Bellpac* requirements.[3] Power and ground are distributed to the circuit pack positions by printed paths. The signal interconnections are in the form of 30 AWG single-ended and twisted-pair wire, terminated by wire-wrap connections.

3.7 Power

Duplicated 140-volt and 24-volt power feeds are required by the DIF. The 140-volt power is used for powering the power units which provide ±5 volt and +12 volt power for the integrated circuits. The 24-volt power is used for the alarm relays, power control circuitry, IPUB drivers, protection switch relays, and indicator lights. The DIF draws 8.3 amps from each 140-volt bus and 1.7 amps from each 24-volt bus.

3.8 Operating environment

The DIF was designed to operate in the No. 4 ESS office environment. This required that it be capable of operating over a temperature range of 4°C to 38°C and a relative humidity range of 20 percent to 55 percent, except for a total of 15 days per year during which time it must operate from 2°C to 50°C and 20 percent to 80 percent relative humidity for no more than 3 days at a time. To achieve these objectives, hermetically sealed integrated circuit (IC) packages or beam-

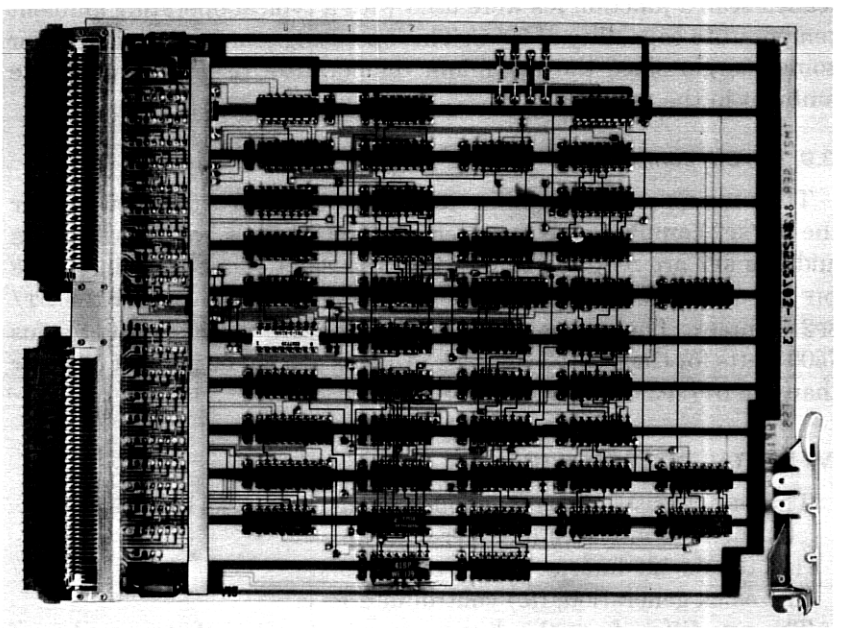

Fig. 9—DIC 4-layer PWB.

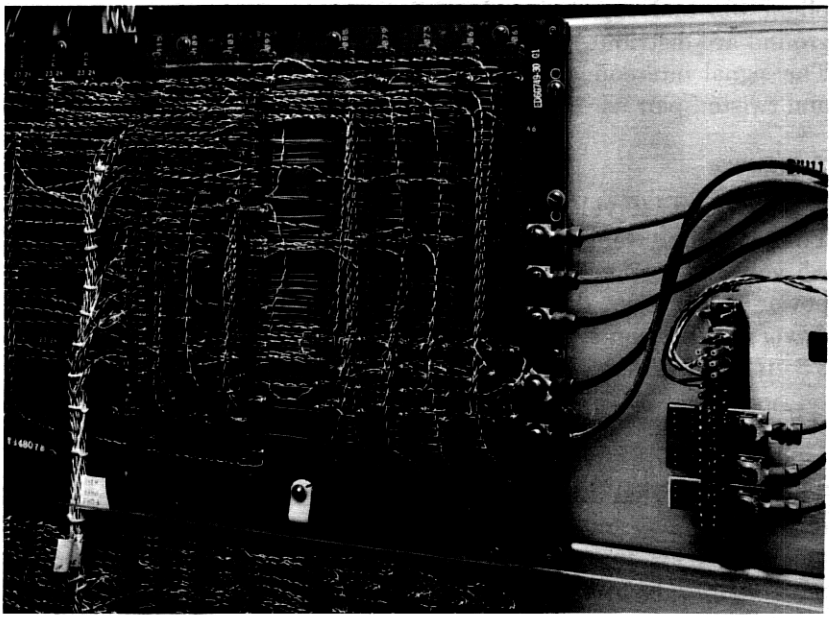

Fig. 10—DIU backplane.

leaded sealed junction ICs were used with a typical operating ambient temperature capability of 90°C. This operating temperature allows for some margin since the maximum temperature rise above the aisle ambient in the DIF is 30°C.

3.9 DT/DIF comparison

The DIF represents a significant size, power, and cost reduction over the DT/SP2 complex. One DIF interfaces 160 T1 lines, whereas four DTs and one SP2 are required to interface this same number of lines. The DIF lineup is 9 feet 9 inches long versus 23 feet 10 inches for the DT/SP2 complex. The power dissipation of the DIF is 2500 watts versus 7500 watts for the DT/SP2 complex. A DIF is around 50 percent cheaper than the DT/SP2 complex.

IV. DIGITAL INTERFACE UNIT

The DIU consists of the four following major blocks (see Fig. 11):

(*i*) the DS-1 interface (one per-digroup),
(*ii*) the DS-120 interface,
(*iii*) clock generation, (*iv*) control interface.

The DIU differs from the digroup terminal unit in that a number of per-digroup functions which were performed using common control

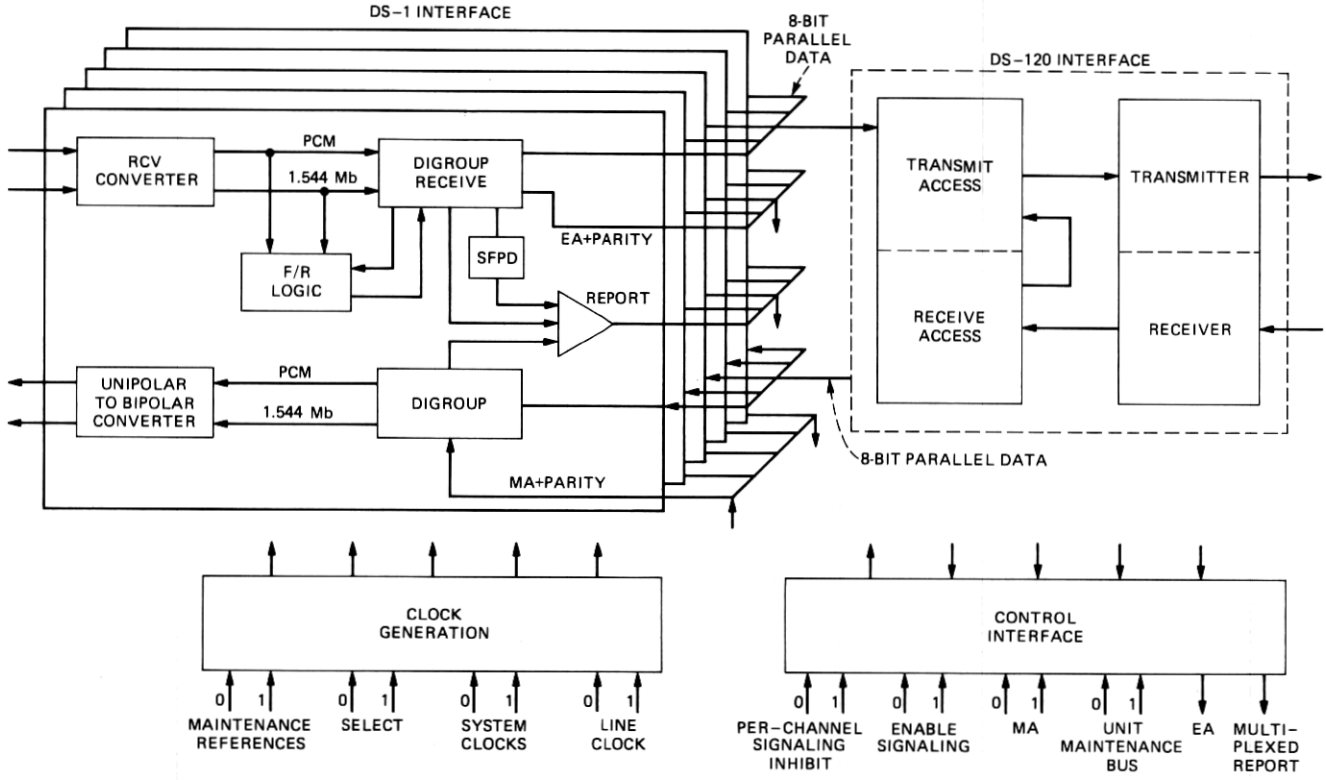

Fig. 11—DIU block diagram.

techniques are now done on a per-digroup basis. This permits the use of custom LSI, which can be shared with other systems and reduces the complexity of maintenance as test vector generators and references are no longer required.

4.1 DS-1 interface

Figure 12 is a detailed diagram of the DS-1 interface which consists of a receive section, a transmit section, and a maintenance and report section.

4.1.1 Digroup receive function

It is the function of the DS-1 receive circuitry to terminate the T1 line, to recover line timing and data, and to convert the incoming data to a form suitable for multiplexing with other digroups into a DS-120 level.

The conversion from bipolar to unipolar, clock extraction and signal regeneration are done with a custom LSI receive converter complementary bipolar integrated circuit chip. The output of this chip is a single-rail unipolar PCM serial stream and a properly phased 1.544-Mb clock. These signals feed the framing and receive logic chip (F/R) and the digroup receive chip (RCV LSI) both of which are custom LSI devices. The F/R recovers the DS-1 data and signaling framing. To function, this chip requires an external line channel counter which is located on the RCV LSI (see Fig. 13). Framing is accomplished by the F/R interrupting the line channel counter a sufficient number of times to restore framing.

The F/R chip supplies a signal indicating a framing error. This signal is used to estimate the error rate and to block the updating of signaling information. The chip also generates two signals which mark the A and B signaling frame.

The major per-digroup receive functions are contained in the RCV LSI (Fig. 13). The primary function of the RCV LSI is to reduce the incoming PCM data 125-μs frame to a specific 23.4-μs digroup interval, to recover A and B signaling bits, and to compensate for differences in line-frame frequency and phase with respect to office-frame frequency and phase. Although the RCV LSI recovers both A and B channels, only the A channel is implemented in the present version of the DIF.

PCM data is converted from serial to parallel (S/P) form and complete 125-μs frames are stored in an A and B random-access memory (RAM) store. (A and B here should not be confused with A and B signaling bits.) The stores are alternately written and read so that in general, while one store (e.g., A) is being written under line timing, the other store (e.g., B) is being read under office timing. If the relationship of line frame to office frame allows reading and writing the store in such

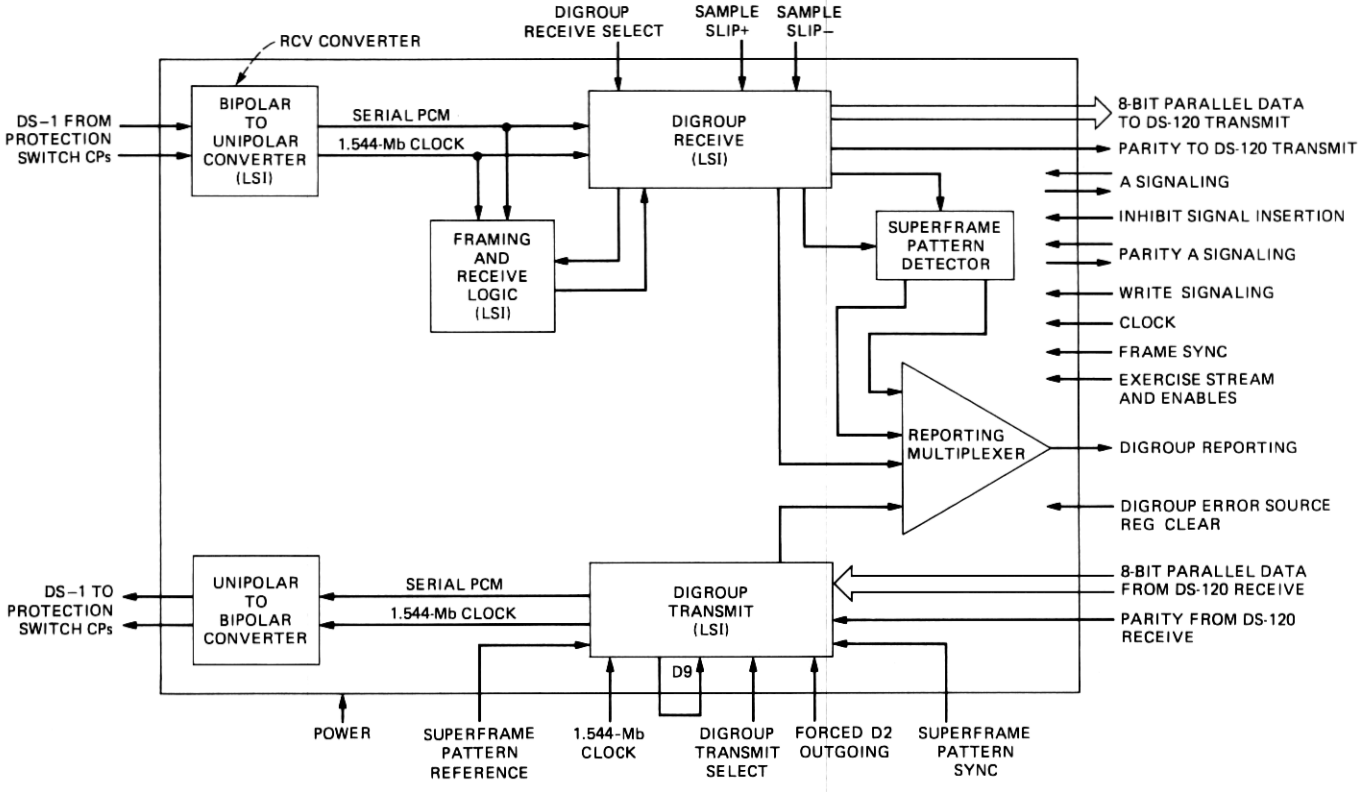

Fig. 12—DS-1 interface.

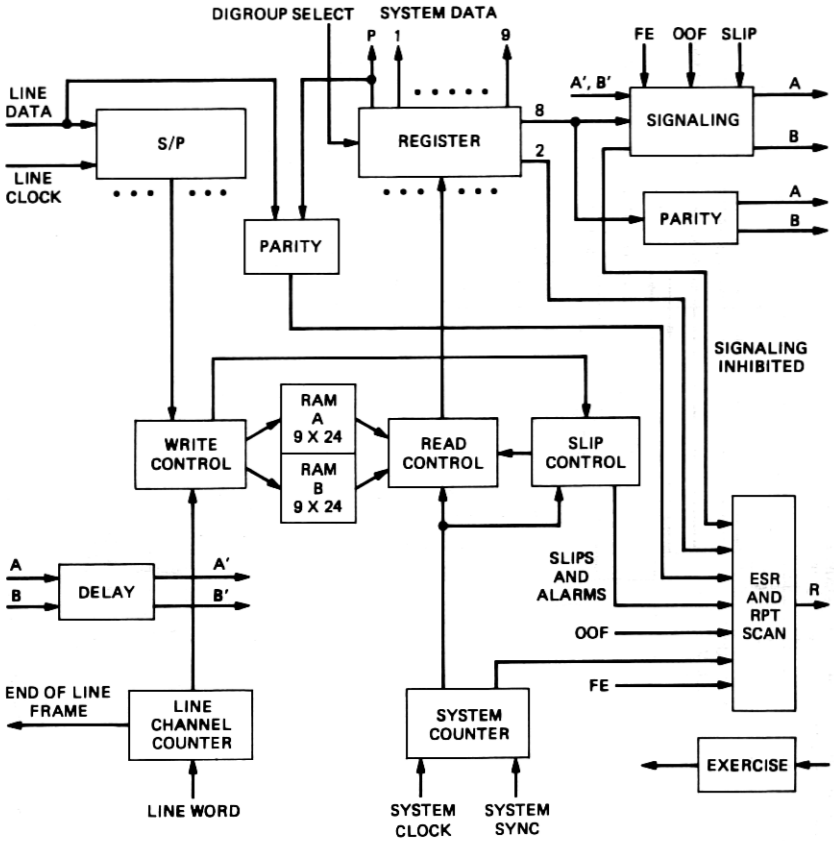

Fig. 13—RCV LSI block diagram.

a manner as to intermix frames, a slip control generates a double A read to correct the situation.

PCM data through the stores is maintained by serial parity over a frame of data. Parity generated at the input to the RCV LSI is stored and compared with parity generated at the output of the A and B stores. In addition to parity, the framing bit, D9, is passed through the stores for maintenance purposes. This is possible since the contents of the D9 bit are defined by signals from the F/R chip.

Extraction of signaling information occurs at the output of the A and B RAMs. Signaling information is stored in two 24-bit shift registers (signaling) and the data is checked by parity. Data is entered into these stores under the command of signals from the F/R chip which indicate when to extract signaling information.

The RCV LSI also contains a detector which determines if PCM bit 2 is held at zero for all frames in a 32-ms interval. This signal is forwarded

to the DIC [via the error-source register (ESR) and report (RPT) scan], which, in turn, times it for the remote (yellow) alarm.

A number of functions are combined on the RCV LSI chip to form a report function stream:

(*i*) line bit D2 stuck at zero,

(*ii*) out of frame (OOF),

(*iii*) framing error (FE),

(*iv*) positive slip (positive slip occurs when the system frame exceeds the line frame rate),

(*v*) negative slip,

(*vi*) alarms.

These reports are used by the DIC to determine the status of the receive portion of the DS-1 interface.

The RCV LSI contains a number of matchers and alarms. Detected failures are combined into a single alarm list. To determine if the alarms and matches are functioning, exercises can be sent to the RCV LSI, which generates alarms without interfering with the processing of data. A digroup clear input is used to clear alarms.

4.1.2 Superframe Pattern Detector

The Superframe Pattern Detector (SFPD) shown in Fig. 12 maintains the integrity of the data flow through the RCV LSI by checking that the D9 bit contains the subframe pattern and that the phasing of this pattern corresponds to that derived from the F/R chip. To avoid the complication of accounting for slips, framing errors and out-of-frames, the SFPD is inhibited during these states.

The ability of the SFPD to detect superframe pattern errors is tested by frame-resident exercise functions (Sections 4.5 and 6.5).

4.1.3 Digroup transmit functions

In the transmit direction, the DS-1 interface receives parallel data, plus even parity from the DS-120 interface. The main functions of this portion of the interface are to select the data for the appropriate digroup, convert it to a 1.544-Mb/s serial stream with the appropriate framing information, and to insert signaling information at the appropriate time. The transmit circuitry consists of two major blocks, the digroup transmit chip (TMT LSI) and the unipolar to bipolar converter, which is comprised of discrete components. The major functions are performed in the TMT LSI (see Fig. 14). Although each digroup interface receives its data during a different 23.4-μs interval, the outgoing DS-1 data are frame and framing pattern aligned for maintenance purposes.

The TMT LSI, like the RCV LSI, contains an address generator, a read-write control, and a RAM divided into A and B sections. While a whole frame of data is being written into one section under system timing,

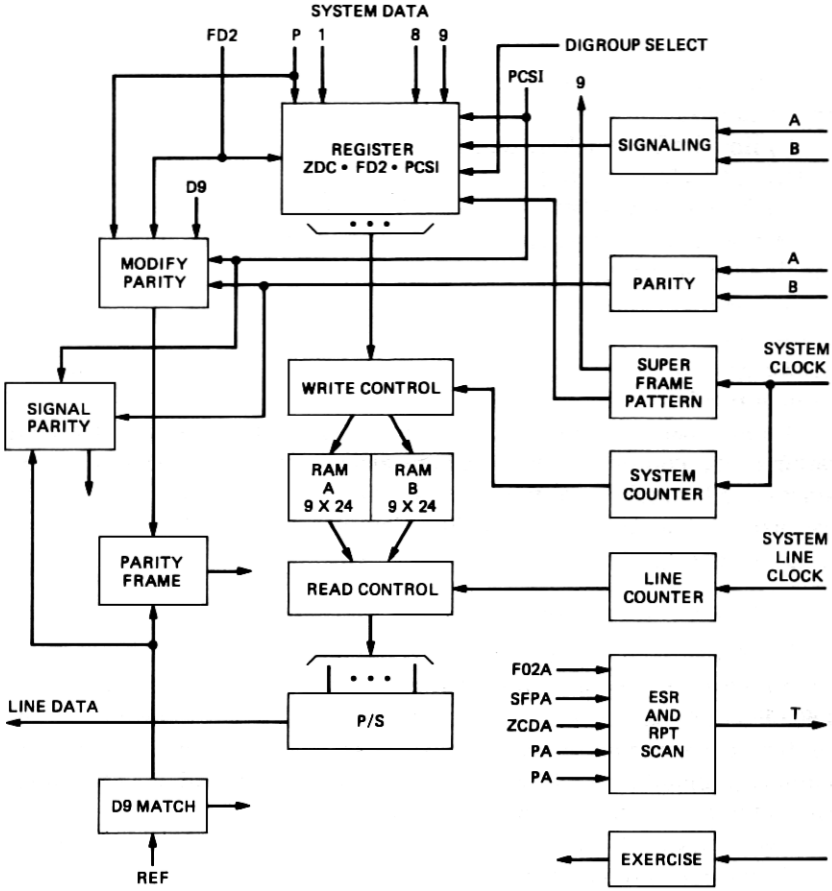

Fig. 14—TMT LSI block diagram.

the other is being read under line timing. Two stores are used for convenience as this allows all digroups to be checked with a reference pattern at the same time on the T1 line side. The TMT LSI also receives signaling information (plus parity) from the control interface (CI) (see Fig. 11) for the A signaling channel, which it stores in a 24-bit shift register in which it is recirculated until an update interval occurs. A superframe pattern generator driven by a synchronizing signal from the CI defines the time at which signaling insertion should take place. The TMT LSI has the capability of handling both A and B signaling storage and insertion, but in the present DIF, only the A channel is used. The TMT LSI also has the ability to prevent the insertion of signaling data in bit position 8 via a control signal from the CI. When signaling is inserted, the parity over data is altered.

At the output of the TMT LSI, serial parity over a frame of data is compared against stored input serial parity over the same data. In addition, the output serial parity over D8 for a frame is compared against stored output serial parity over D8 for signaling insertion maintenance. The TMT LSI contains a zero code detector and a forced D2 inserter. The zero code detector forces bit 7 to a one if all bits of a word are zero. The forced D2 inserter under command of the CI forces bit 2 to a zero to transmit the yellow alarm to the distant terminal. This action is initiated by a command from the DIC. The forced D2 feature is maintained by returning a status bit to the CI, while the zero code detector is duplicated.

As in the RCV LSI, the TMT LSI multiplexes all of its status and alarm bits into a single stream which is forwarded to the reporting multiplexer. In addition, it has exercises to test the various matches and alarms.

The serial PCM data from the TMT LSI is modified by the unipolar to bipolar converter. This block is made of discrete components and puts out a bipolar signal of the proper amplitude and shape for driving a T1 line.

4.2 DS-120 interface

The DS-120 interface consists of four functional blocks: the receive access, the receiver, the transmit access, and the transmitter (see Fig. 11).

4.2.1 Receiver and receive access

The DS-120 termination contains an analog line receiver similar to the one used in the digroup terminal whose function is to terminate the coaxial cable, amplify the PCM data signal from the TSI, and supply a sampling clock for the data. The data signal framing is determined and the signal is converted from the bit/bit complement serial format of the DS-120 link to an 8-bit parallel form with parity.

The receiver contains a squelch function. If two successive frames have pair violations, the squelch is applied. When the receiver recovers frame, the squelch hangs over for 8 ms. The squelch applies an all-ones code to the outgoing data to prevent analog carrier overload.

In the receive access, test vectors (derived from unit clock signals) are inserted in spare time slots 127 and 0. These vectors cannot propagate through the per-digroup equipment. However, time slots 125, 126, 127, and 0 are forwarded to the transmit access. A two-time slot delay occurs between the transmitter and the receiver. Hence, when a DIU is looped on itself, data in time slots 127 and 0 are sent to 1 and 2, 3 and 4 are sent to 5 and 6, and so forth. Thus, when looped, the vectors eventually occupy all working time slots and can be used

to test the looped DIU. This condition pertains to the spare DIU when not in service and for a DIU when it is protection-switched while in an out-of-service condition. If a DIU cannot be protection-switched, it is possible to notify a distant office of an out-of-service condition under most circumstances, by sending an exercise that blocks the transmission of the framing signal. Time slots 125 and 126 are looped to permit the TSI to test transmission to and from the DIF.

4.2.2 Transmitter and transmit access

In the transmit access, the receive and transmit streams are combined together. To ensure proper operation of the access gates, parity over the spare time slots 125 and 126 is even while transmit data parity is odd. A failure to multiplex generates a parity failure. The conversion of data from a parallel form to the bit-bit prime serial format of the DS-120 link is protected by recomputing serial parity, in test vector time slots 127 and 0, and comparing it with a reference parity for those time slots.

The transmitter consists of a TSI line driver similar to the type used in the DT which amplifies and buffers the DS-120 signal to drive up to 1000 ft of 100-Ω coaxial cable.

4.3 Clock selection, generation, and decoding

To reduce the number of leads between the DIC and the DIUs, each DIU generates its own clock chains. All DIUs must be frequency- and phase-locked to the DIC; this is achieved by furnishing to each DIU six signals over a duplicated link:

(*i*) System signals
  (*a*) 16.384-Mb/s square wave clock,
  (*b*) 8-kHz synchronizing signal,
  (*c*) 31.25-kHz synchronizing signal.

(*ii*) Line signals
  (*a*) 1.544-Mb/s square wave clock,
  (*b*) 666.67-kHz synchronizing signal for the T1 line superframe pattern,
  (*c*) 8-kHz synchronizing signal for the T1 channel counters.

In addition to the above signals, two reference signals are sent to the DIUs to ensure that the local clock generation is proper. The first is a signal representing parity over all the system clocks generated in the DIU having periods from 250 μs to 32 ms. The second is a reference signal which is compared with superframe patterns generated in the digroup transmit chip for maintenance purposes.

All of the above signals are duplicated and selected by a 1A Processor command via the DIC.

4.4 Control interface

Information exchanged between a DIU and the DIC is processed by the CI.

From the DIC to each DIU there are four data streams:

(*i*) M signaling (A channel outgoing),
(*ii*) enable signaling (ENSIG),
(*iii*) per-channel signaling inhibit (PCSINH),
(*iv*) unit maintenance bus (UMB).

These signals are originated in each of the controller halves, distributed to all units serially via balanced drivers on a duplex basis, and selectable by a control command via the CI.

M signaling contains the multiplexed signaling data to be distributed to the various DIUs and digroups. Each DIU has a hardwired time of unit code (TU), so that a particular DIU can only off-load signaling during a unique 125-μs interval every 8 ms. Data integrity is maintained by odd parity in time slot 127. The ability of a DIU to correctly decode the TU is checked every 8 ms by the DIC via the looped spare time slots. The CI generates serial parity over each stream and forwards them to the five DS-1 interfaces along with the multiplexed signaling stream. The per-digroup circuitry picks off its signaling data based on a digroup select clock signal. If a parity failure occurs for the incoming data, the CI alarms but does not block updating as that would require a 128-bit store. Signaling information, however, cannot be updated unless the ENSIG is in the proper state. Enable signaling permits the DIC to block updating of DIUs from a faulty DIC half. If the signaling failure occurs because of a detected DIU failure, the unit is protection-switched.

Per-channel signaling inhibit allows a DIU to block the insertion of signaling information into the DS-1 stream for any channel in that DIU. The CI contains a 128-bit store which holds the PCSINH status. This signal permits the DIU to pass full 8-bit information for common-channel interoffice signaling (CCIS) or for permanently connected special service trunks. The CI loops and multiplexes the received spare PCSINH time slot data with the M signaling spare time slot data so that the DIU can maintain the distribution to the DIU. Distribution is checked every 16 ms by the DIC. Data to the DIUs are protected by odd parity in time slot 127. The CI distributes PCSINH on a looped basis to all digroups and the parity over the data returned to the CI is compared against the parity received from the DIC.

From each DIU to the DIC two streams of data exist:

(*i*) multiplexed report stream,
(*ii*) E signaling stream (A channel incoming).

The report stream is composed of the per digroup reports (24 bits per digroup) as shown in Fig. 15. Thus, the report stream is composed of 120 bits for each DIU. The various per-digroup report streams are multiplexed in the DIU backplane and this multiplexed stream is then combined with the 8 common alarm bits in the CI. The DIU repeats the reports every 125 μs. The common control reports are autonomously cleared every 32 ms, whereas the per digroup reports are cleared via a command from the DIC. The autonomous clear for each DIU occurs at staggered times so that the DIC can interrogate the DIUs without danger of missing a failure report because of a DIU self-clearing action.

The multiplexed report stream which is transmitted in simplex form to the DIC from each DIU is driven from an unbalanced source and its ability to transmit data is maintained by frequent controller exercises.

The signaling stream is comprised of five multiplexed digroup E signaling channels and looped channels. Seven spare signaling channels (121 to 127) received by the CI are looped back to the DIC via the A signaling stream with a 125-μs delay. The DIC will only look at the

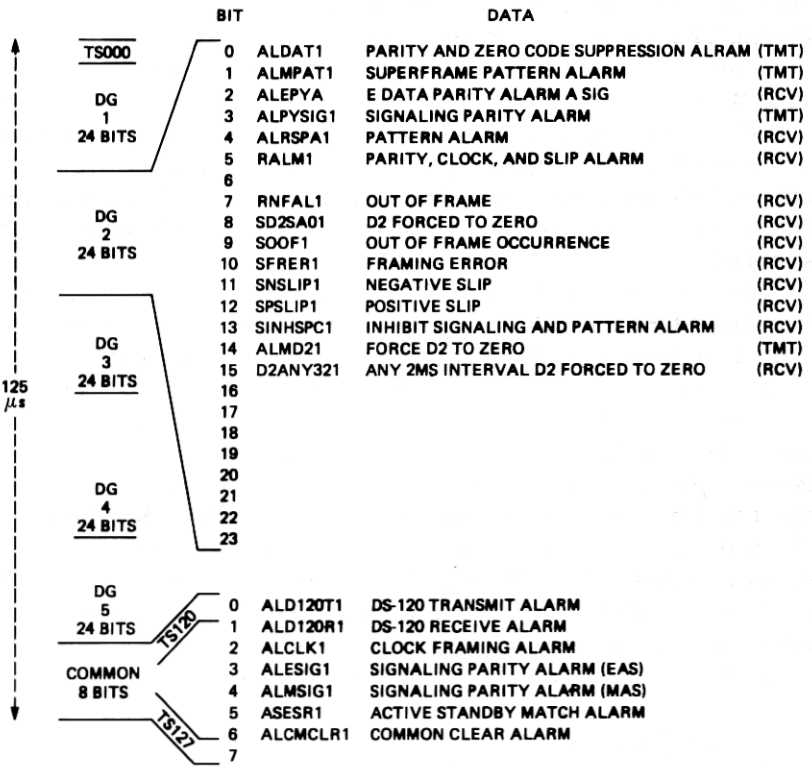

Fig. 15—DIU multiplexed report stream.

looped data during its normal E scan once every 8 ms. Reception of M signaling and PCSINH is tested by verifying that the looped channels contain these signals during alternate 8-ms scans. The signaling stream, like the report stream, is simplex from each DIU and driven by an unbalanced source.

The UMB is the communication link from the DIC to the DIUs. It is via this stream that the DIC can exercise the DIUs and insert the remote alarms. The UMB line format is shown in Fig. 16. Data on the maintenance bus is frame synchronized to DIU clocks but is not distributed by decoding a TU. Instead, a DIU decodes six address bits, U0 to U5, to determine if the data is for it. A seventh bit, ALLUN, is used to address all DIUs simultaneously.

The function enable determines the action to be taken by a DIU. If the action to be taken is for a specific digroup, then the digroup identity will be flagged. The exercise fields are not coded, i.e., each bit of an exercise word is a specific common or digroup exercise state.

The UMB distribution to the DIUs is maintained by odd parity over all data in TS 127 within the DIC.

4.5 Exercises

The DIU contains a more comprehensive list of exercise routines than the digroup terminal unit does. Normally, exercises are used to test on a periodic basis the error source registers and matchers in the DIU. However, in the DIU a set of exercises tests the per-digroup equipment when the DIU is protection-switched. More specifically, the DIC has a set of diagnostic routines that tests for the ability to:

(*i*) detect framing errors,
(*ii*) reframe,
(*iii*) detect forced D2,
(*iv*) send forced D2,
(*v*) transmit and receive signaling.

For items (*ii*) and (*iii*), time limits are set on the response of the DIU so that deterioration of the F/R logic to respond properly can be detected.

V. DIGITAL INTERFACE CONTROLLER

5.1 Design objectives

The DIC represents the latest development in the evolution of the No. 4 ESS peripheral frame controllers.[4] Its design draws heavily on experience accrued over earlier transmission/switching designs, and capitalizes on microprocessor technology to allow flexibility to respond to change.

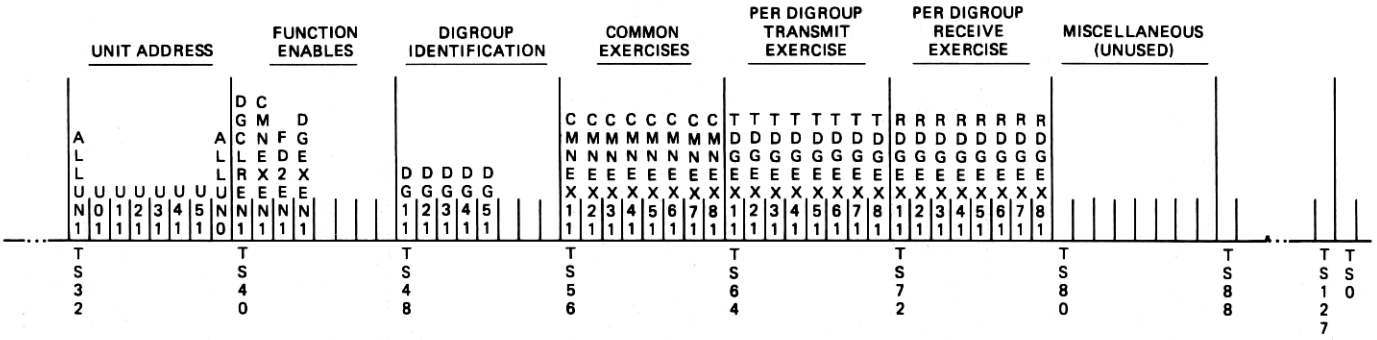

Fig. 16—DIU (DIC maintenance bus format).

Functionally, the DIC supports the DIUs in much the same way that the digroup terminal controllers and the related portions of the Signal Processor 2 supported the digroup terminal units. It provides for the collection, processing and distribution of supervisory and address signaling information, a reliable clock source, and an interface to the PUB of No. 4 ESS.

The DIC provides reconfiguration of the DIUs under fault conditions by controlling the protection switching of the spare units and appropriately redirecting signaling interchanges. The transmission facility maintenance and DIU maintenance are provided via a maintenance microcomputer developed in the controllers, around the *Bellmac**-8 microprocessor.

Many software considerations went into the design of the DIC. The repertoire of operational (call processing related) peripheral bus orders was chosen to be compatible with those of the Signal Processor 2. This minimized the associated call processing software development. The maintenance software in the 4E5 generic was dramatically restructured to allow for "intelligent" controllers and to minimize and modularize hardware dependent software. The DIC was designed to minimize generic DIU reconfiguration software, and controlled most DIU reconfiguration actions with frame resident "firmware."

The DIC was designed so that a significant portion of the controller hardware and diagnostic software could be used in common with the Peripheral Unit Controller (PUC) associated with the mass announcement system feature in 4E5.[5] This minimized overall development effort, and allowed the PC to capitalize on the larger scale of manufacture of the DIC. The maintenance software development for the PC was also facilitated, in that common hardware characteristics minimized differences in the frame-dependent software.

5.2 Controller architecture

Figure 17 shows the overall architecture of the DIC. Since a failure of a DIC could affect up to 3840 trunks, the controller is fully duplicated. Either controller can support all the DIUs, while its mate is being diagnosed and repaired. The two controllers are independently powered and are provided independent clock inputs (Master Timing Links or MTLs) from the associated TSIs or echo suppressor terminals. Two MTLs are provided per controller.

The controllers derive their internal timing from either connecting MTL, and each controller has its own independent countdown chain. Synchronization signals are provided by both controllers to the DIUs which select one and only one controller's timing to drive the internal

* Trademark of Western Electric.

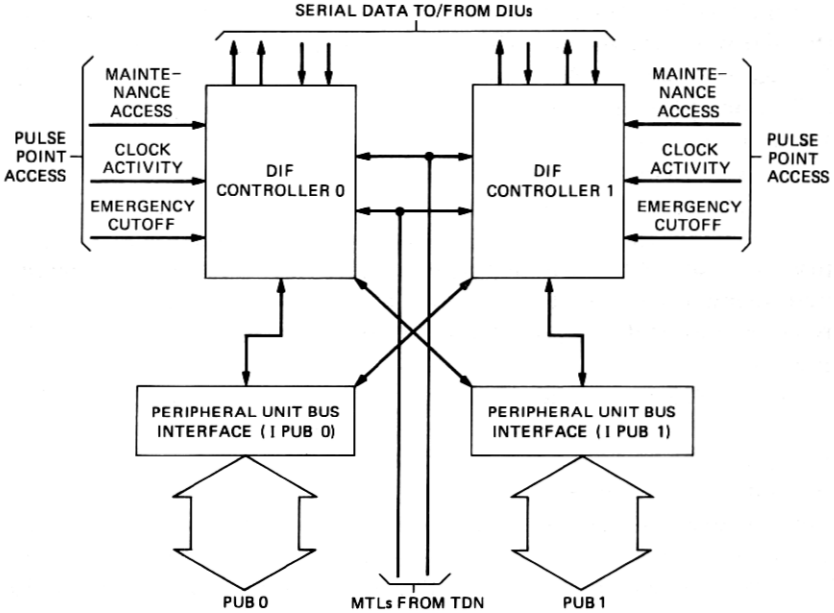

Fig. 17—Duplex DIC overview.

DIU functions. The ten signals provided to the DIUs allow their DS-120 outputs to be synchronous to the No. 4 ESS time division network and also determine the outgoing T1 line frequency of 1.544 Mb/s.

The PUB interface, which is also duplicated, terminates in the bus access circuits as shown in Fig. 17. This interface is fully configurable so that either controller can take its input from either bus, and can reply on either or both buses. In addition to the PUB, the 1A Processor is able to access the DICs via pulse points. These pulse points are employed for recovery actions, such as putting a controller in the maintenance state, selecting clocks, or disabling a controller's protection-switch capability. The pulse points are provided via independent signal processors for each controller.

The maintenance and operational data which pass between the DICs and DIUs are time-division multiplexed. Consequently, a transmission path is not required from each DIU for each signal. This results in simplified intraframe cabling and minimal select circuitry in the DIUs.

5.3 Interface to the peripheral unit bus

The DIC terminates the entire PUB. The bus consists of four duplicated bus groups:

- The PU enable/address bus, which conveys address information that determines which frame in the No. 4 periphery should respond to a particular peripheral order.

- The PU write bus (PUWB), which conveys the data to be accepted by the DIC and processed.
- The PU reply bus (PURB), which is used to return data to the 1A Processor.
- The PU control bus over which control and maintenance information is transmitted to and from the peripherals.

The bus access circuitry terminating PUB0 is independently powered from that terminating PUB1, and both are powered independently of either controller. In this way, a failure of either bus access circuit or either controller does not affect the balance of the controller and bus circuitry.

As part of the bus access circuitry, "bus clamps" are provided to prevent a faulty controller from "babbling" onto either bus. These clamps are controlled with a combination of the maintenance access pulse point and either the member interrogate or group interrogate bits of the control bus. When a bus or controller is powered down, the clamps are manipulated by power sequencing logic to prevent babbling.

The PUB signals incoming to the DIC are terminated in series receivers before passing on to the next peripheral frame. So that replacement of a receive pack does not interrupt the continuity of these signals, bypass resistors are provided. If the DIF is the last peripheral on the bus, optional terminating networks are employed.

5.4 Simplex controller architecture

Figure 18 shows the internal architecture of one of the duplicated controller circuits, commonly referred to as a simplex controller. PUB orders coming to the DIC are stored in the receive logic, and tested for validity. If the order is directed to this particular DIF and obeys the appropriate protocol, the receive logic interrupts the Executive Controller (EXEC). The EXEC routes the order to the appropriate function via the internal bus. If the order requires a reply, the data from the subject function is routed through the internal bus and into the reply logic. The reply logic then controls the appearance of the reply data on the PURB.

The function of the receive logic is to store the bits received from the PUB and test the received information for the format and addressing. In addition, it provides synchronization between the asynchronous PUWB data and the synchronous interface of the EXEC controller. The reply logic provides the necessary reply formating of the information returned to the 1A Processor via the PURB. There are a number of special reply bit fields (e.g., maintenance data) that are required of the DICs. The receive and reply logic provides the control signals to handle these special requests from the 1A Processor.

The internal bus structure of the DIC consists of a multiplexed bus. This bus structure was chosen to minimize fault susceptibility, to

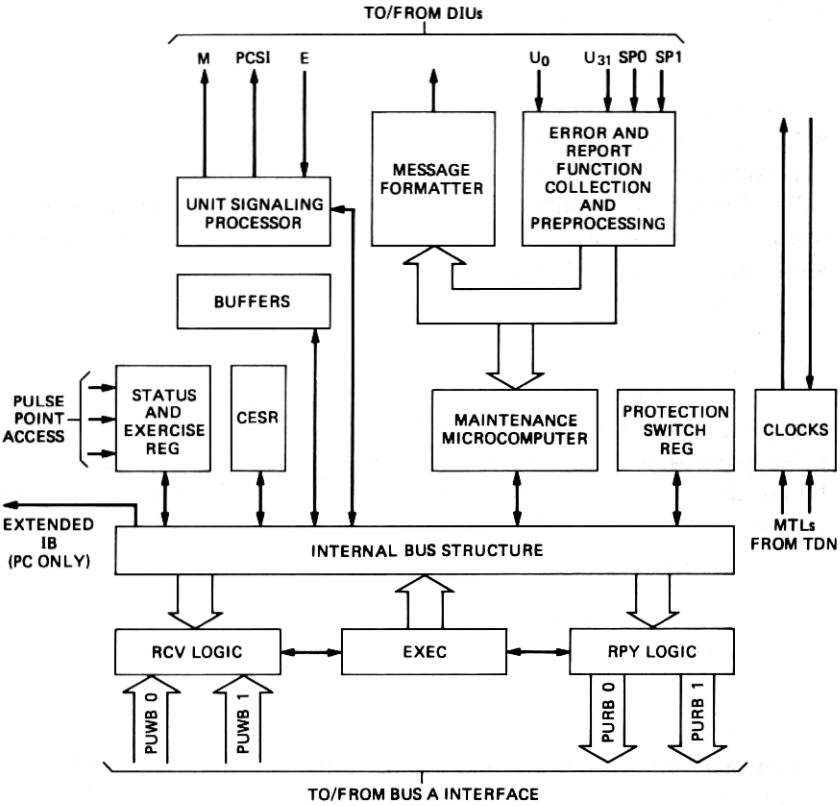

Fig. 18—Simplex DIC.

alleviate the necessity for multiple high-powered drivers, and to allow better control of the bus topology. All internal bus lines are looped back to the internal bus circuit pack driving them, for impedance termination and maintenance. The internal bus, which is under control of the EXEC, selects one of 16 possible ports or data sources to be placed onto the internal bus. Each port consists of 24 data bits and a parity bit. Each functional entity within the DIC is considered a port on the internal bus. The receive and reply logic is treated the same as the other functional entities within the DIC. This parallel high-speed structure allows flexible communication within the DIC.

The EXEC configures to the internal bus multiplexers with four bits of select information to specify which data of the 16 ports it will route through the internal bus. These four internal bus select leads are duplicated for maintenance. The EXEC also sends to the functions within the DIF six bits of source address, six bits of destination address, four bits of internal bus operation code, and an internal bus load pulse. The source address specifies which function will be placing the data

onto the internal bus. This allows for premultiplexing, where necessary, in the various functions. The six-bit destination address specifies the recipient of the data. The function addressed loads its internal register with the data on the internal bus upon receiving the load pulse. Each function also returns to the EXEC a source and destination acknowledge. This is an acknowledgment of a function's receipt of a valid source or destination for one of its registers. This acknowledge allows the EXEC to check the source and destination decoders in the various DIF functions.

The controller handles all internal communications in the DIC and communication with the 1A Processor. The EXEC is a bit-sliced, bipolar microprocessor. The microstore is modular with a maximum capacity of 4000 words of microprogram. The EXEC contains a 16-level priority interrupt control circuit, a 12-bit microsequencer, and an 8-bit arithmetic logic unit (ALU). Internal bus access logic allows the ALU to access 24-bit data from various DIC functions. Inputs to the EXEC consist of interrupt signals from various functions within the DIC. These signals indicate a request for a particular task routine to be executed.

The microprogram of the EXEC consists primarily of an interrupt handler and numerous special purpose task routines to control internal bus transactions. Interrupts from the receive logic (i.e., from the 1A Processor), from the maintenance microcomputer, from the unit signaling processor, and from a real-time clock (10 ms) source can all initiate such task routines. As an example of the microprogram execution, consider the receipt of an interrupt signal from the receive logic. The execution of an EXEC order can be broken down into three parts. The first of these three is the time from the presentation of the interrupt source until the interrupt is checked and detected. This is a noninterruptible state and consists of tasks that cannot be subdivided or "segmented." An example of a noninterruptible task is a bus transfer with an internal bus client with volatile data.

The second part of EXEC order execution is the time to transfer control to the interrupting routine. This is the time for the interrupt to be acknowledged, for control to be passed through the interrupt jump vector table, and for the resetting of the interrupt circuit. The third portion consists of the actual execution time of the requested function. The EXEC must complete the data transaction within a finite time window or the receive logic triggers a controller alarm. This is done as a sanity check against the EXEC.

In addition to interrupt initiated task routines, the EXEC contains microdiagnostics which are invoked by 1A Processor resident diagnostic programs. These microdiagnostics extensively check the ALUs, microsequence controller, and microprogram store of the EXEC.

The overall configuration of the DIC is determined by the data in the controller status register. Routing of PUB data, MTL selection, and controller selection are determined by the state of this register. The exercise register is used during diagnostics to exercise the fault-detection circuitry of the DIC and verify their ability to report. Under normal conditions, all hardware fault detector outputs are stored in the controller's error-source registers. Each major functional element in the controller has a local error-source register, and a summary of these are recorded in the primary controller error-source register (CESR). This register is interrogated by maintenance software during a fault condition to isolate the fault to a particular controller. Unlike some of its predecessors, a minimum of fault detection in the DIC relies upon cross controller matching. Many processing elements are duplicated within a simplex controller to assure autonomous maintenance even under simplex operation.

The disposition of the spare DIUs is controlled by the protection switch registers. The protection switch register outputs of the two controllers are logically ORed in the protection-switch equipment. Special cutoffs are provided to prevent a faulty controller from affecting a protection switch. The power sequencing circuitry activates these cutoffs (in addition to software control) when a controller is powered down.

The operational functions of E signaling reception, M signaling distribution, and PCSINH are controlled by the unit signaling processor. The unit signaling processor function includes the collection and distribution of supervisory signaling, as well as dial pulse reception and outpulsing. The signaling processing is determined by state translation firmware, and therefore can be modified to respond to changes in pulse width requirements or general timing changes. The signaling on all trunks is processed every 10 ms, as initiated by the EXEC in response to the real-time clock interrupt.

The results of signaling processing are reports which must be communicated to the 1A Processor. The SP deposits reports via the internal bus in one of four "buffers" or scratchpad RAM regions. These buffers are designated high priority, low priority, seizure, and digit buffers. The EXEC routes the data into the buffers, and administrates the appropriate read and write pointers.

The 1A Processor periodically polls the DIF, just like an SP2, to determine if any reports are present. If reports are present, the DIF acknowledges the poll, and call-processing software reads the buffers and either directs the unit signaling processor to continue processing that trunk, or proceeds to connect a path through the No. 4 ESS time-division network.

The DIC has a maintenance microcomputer (MMC) to maintain the

DIUs, aid in fault-recovery actions, or participate in diagnostics of other portions of the controller. The MMC has a byte-addressable serial bus with which it sends information to the DIUs. The DIU report streams are converted to parallel format and deposited in the MMC RAM via direct memory access (DMA). Following a complete DMA cycle, which occurs once every 32 ms, the preprocessing hardware generates an interrupt which initiates processing of the DIU alarm data. In this way, DIU common alarms are "hit timed," local and remote (T1) alarm indications are detected, and facility reports such as slip, out of frames, and error rates are prepared.

The MMC is built around a *Bellmac*-8 central processing unit (CPU) which is duplicated within a simplex controller for fault detection. The MMC executes tasks under direction of an interrupt-driven operating system called "OS8" written in the C programming language. All of the MMC resident application programs are written in C as well, with the exception of real-time intensive unit alarm processing tasks which are programmed in assembly language.

The 1A Processor can access any memory location in the MMC via DMA. Autonomous reports from the MMC are deposited in a portion of buffer RAM designated the maintenance buffer. Like the operational buffers, this is administrated by the EXEC and periodically polled by the 1A Processor.

The MMC is capable of (macro) peripheral order expansion, under direction of the 1A Processor. In this way, complex tasks such as initialization and unit reconfiguration are relegated to the MMC and the 1A Processor maintenance software is considerably simplified. This was consistent with the restructuring of maintenance software, a significant part of the 4E5 generic.

Each simplex DIC independently derives clocks that drive the controller's circuitry, and which are sent to the DIUs. The clocks are derived from the MTLs originating in the time-division network (TDN) of No. 4 ESS. In this way the DS-120 signals of the DIUs are fully synchronous with the timing of the associated TSI.

To preserve commonality between the DIC and the PC, two major features were a part of the design. First, a port of the internal bus was equipped with bidirectional drivers/receivers to form the extended internal bus. Since this extended bus is tristate, and any client output shorting low would affect all clients, special extended bus cutoff signals are provided. These can be used to quarantine any extended bus client and allow the PC to continue operations.

The PC does not have DIUs, so all DIU-related hardware was packaged in the lower shelf of the DIC. By deleting this shelf and the associated circuit packs, the PC did not bear this unnecessary cost burden and 80 percent of the DIC hardware is still applicable to the PC.

VI. MAINTENANCE AND RECOVERY STRATEGY

6.1 Overview

The maintenance and recovery strategy employed by No. 4 ESS is structured into five levels. These levels are phase level, interrupt level, interject level, base level maintenance, and diagnostic isolation. Of these levels, the first four are related to identifying and isolating the failing subunit and recovering the system. The fifth level attempts to isolate the failures within a frame and is used to maintain the operational soundness of the frame. As the service impact of a failure increases, the level of response to that failure is altered. Within this system, each frame can specify the initial level of recovery to be associated with a failure.

With the introduction of the DIF, an additional level of maintenance was added. The new level, routine exercise, is used to detect faulty DIUs before they have been able to adversely affect the service they provide. In the next few sections we present a description of the maintenance and recovery strategies specifically used to support the DIF. Table I shows the maintenance and recovery strategies to be covered.

6.2 Phase level

Phase level is the severest of the recovery strategies, in that service interruption may be experienced during its execution. Phases are stimulated when a failure seriously hinders the effective operation of the system or if reinitialization of part or all of the system is required for recovery. Phase level is further segmented into four levels. The most drastic, level 4 (phase 4), is the level of phase recovery we will concentrate upon.

Table I—Maintenance and recovery software hierarchy

Phase Level
System Failures Action—reinitialize system
Interrupt Level
Controller (DIC) errors Action—DIC configuration
Interject Level
Unit (DIU) errors Action—DIU configuration
Base Level
Normal operating level Action • Base level maintenance • Diagnostics

A phase 4 cleanly reinitializes the entire No. 4 ESS office, leaving failing frames out-of-service. In previous generic issues, frame initialization during level 4 executed in a strictly serial manner and resulted in simplex operation. In the case of the DIF, a new approach to level 4 initialization was implemented. The DIFs within an office are initialized to duplex operation with the initialization of all DIFs being done in parallel.

To a large degree, this change in initialization design can be credited to the incorporation of a maintenance microcomputer (MMC) within the DIF. The MMC is responsible for the bulk of the frame initialization function. The MMC accepts orders from the 1A Processor, which it expands and executes in addition to executing its internally generated maintenance actions. The internal expansion of 1A Processor orders by the MMC allows the 1A Processor to issue a command to the DIF and have an entire function executed, while the 1A Processor does something else. Thus, during phase 4, a small amount of configuration is performed on the first DIF, followed by the issuance of a frame initialization order. As the frame "init" order is processed by the first DIF's MMC, the 1A Processor directs its attention to the next DIF in the office. This sequence is repeated until all DIFs have been initialized. Upon the completion of the frame "init" order, the MMC returns a success or failure indication. The 1A Processor uses this information to determine the resultant configuration of each DIF (duplex, simplex, or duplex failed). The final result of the phase 4 is a stable operating office.

6.3 Interrupt level

In developing the DIF, it was concluded that errors within the DIC should be reported separately from those errors associated with DIUs to simplify the resolution of faults. It was also decided that DIC errors should be reported at a higher level to prevent the possible loss of all the DIUs in the alarming frame. The F-level interrupt is the mechanism used by the DIF to notify the 1A Processor of failures associated specifically with its controller circuitry. It is the initial level of recovery associated with a DIF controller failure.

Let us suppose that a No. 4 ESS office is operating stably. Further, suppose that a DIC in that office experiences a problem. This problem may be due to a hard failure, the loss of a circuit pack, or it may be caused by a transient failure condition. In any case, the operation of the DIF has been disrupted and needs to be corrected. An interrupt (F-level) is the mechanism used to inform the 1A Processor of the problem and request action to resolve the failure.

A failure detected within a DIC results in the setting of an ESR bit(s) associated with the failed circuitry. The setting of this bit(s) stimulates

the F-level interrupt. The recognition of this interrupt at the 1A Processor invokes the DIF interrupt recovery package, DIFRINTR. It is the responsibility of DIFRINTR to determine the appropriate frame recovery action based upon the error-source signature, the initial configuration of the frame, and the relative frequency of reported ESR bits from the frame.

The actions available to DIFRINTR are basically three. DIFRINTR can decide that the appropriate recovery action is the restoral of the alarming controller. This may be dictated by the transient nature of the failure or the simplex configuration of the frame. DIFRINTR can request the listen-only removal of the alarming controller followed by a diagnostic. The majority of the failures occurring in duplex frames are handled using this option. Removing a controller to listen-only allows it to be isolated from its mate, yet kept entirely up to date until the diagnostic begins. The listen-only removal request allows DIFRINTR a second chance if the wrong decision as to which controller to remove was made. The third option available to DIFRINTR is the zero-start restoral of the mate. This option is executed as a final attempt to preserve the operation of the frame. It terminates all stable and transient cells that are being handled by the alarming DIF. The alarming controller is removed from service while the mate controller is completely initialized. Failure of this option results in the duplex failure of the alarming DIF.

6.4 Interject level and base level maintenance

Both the interject and base level maintenance (BLM) reporting mechanisms are additional methods used to inform the 1A Processor of a problem and request action to resolve the failure. The particular levels of recovery in relation to the DIF are strictly reserved for DIU-associated failures. Within the DIF, these failures are monitored and reported by internal frame processes.

The maintenance of each DIU is performed internal to the DIF frame within its MMC complex. On a 32-ms cycle, data associated with the performance and health of each DIU in the frame is written into the MMC memory spectrum. The MMC real-time DIU maintenance firmware scans this data in search of errors and reports both common alarms and digroup alarms to the 1A Processor.

Common alarms indicate a malfunction affecting the operation of 120 trunks, an entire DIU. To preserve the operation of this alarming subunit, these failures need to be detected quickly. For this reason, they are scanned for every 32 ms. The discovery of a common alarm places that DIU on the hit timing list. If the alarm exists for three consecutive scans (96 ms) of the DIU data, it is classified as a hard failure and reported to the 1A Processor. Common alarms are reported

via the Autonomous Peripheral Unit Trouble bit, in the frame's primary ESR. This causes interject processing to be scheduled by the 1A Processor. Reporting common alarms via an interject is rapid enough (served within 20 ms) to allow the recovery software the option of protection switching a spare DIU for the alarming DIU. Protection-switching preserves stable calls active in the alarming DIU.

Digroup alarms are less critical with respect to trunk effect than common alarms. These failures affect the operation of only a single digroup (24 trunks) within a single DIU. These indicators are scanned once every 384 ms. The lower-scanning frequency for these failures results from the requirement that all DIUs be checked for common alarms every 32-ms cycle. The time remaining in each cycle after this checking permits the scanning of three DIUs for digroup failures. Even with the longer scan time, digroup failures are hit timed for three counts before being labeled hard. Hard failures of the digroup variety are reported via a BLM. A BLM is less system-affecting than an interject and is the result of a failure report being retrieved from the MMC maintenance buffer on base level. The BLM removes the failing digroup from operation.

6.5 Diagnostics

Once the recovery software (outlined above) has isolated an alarming DIC or DIU, the DIF diagnostic is called upon to further resolve the failures. It is the responsibility of the diagnostic to specifically determine which circuitry, if any, has failed. Toward this end, the DIF diagnostics are structured into phases. Each phase is responsible for determining the operational fitness of a specific section of the frame's circuitry. Thus, the operational fitness of the entire frame is based on an all-tests passed condition being achieved by all phases of the DIF diagnostic. A diagnostic phase should not be confused with the system phases discussed earlier.

The DIF controller diagnostic, for example, consists of 23 distinct phases. The phases of the DIF diagnostic are ordered such that the fitness of the frame is checked using an "onion peeling" philosophy. In peeling an onion you start at the outside and work your way inward. The DIF diagnostics are implemented in much the same way. The early phases of the DIC diagnostic start at the PUB and clock interfaces to the frame and progress inward. The early phases confirm the 1A Processor's ability to gain access to the frame. Later phases attempt to verify the functional integrity of the executive processor and check the accessibility and operation of the DIC's internal bus. At this point in the diagnostic, the front end and the bus to the workings of the controller have been verified. The remaining phases of the diagnostic determine the soundness of the maintenance microcomputer complex

and the signal processor complex which are accessible via the internal bus.

The end result of a diagnostic analysis of the controller or subunit circuitry will be the isolation of the circuit pack(s) that have failed. Upon replacement of the faulty equipment, the diagnostic is executed again to ensure that the problem has been corrected before returning the DIC or DIU to service.

6.6 Routine exercise

In addition to the other functions detailed previously, the inclusion of the MMC complex has afforded the DIF the capability of executing tasks (functions) on a routine basis. When the MMC has no requested tasks to execute, it sequentially executes the tasks that have been identified as routine. To date, two tasks, DIU exercising and auditing of the protection-switch state, are run routinely within the DIF.

The DIU exercise routine task uses test vectors to functionally test the operation of each DIU. If a DIU does not respond properly to the test sequence, the failure is reported via BLM. The protection-switch state audit routine task checks for the proper configuration of the spare units. A failure of the audit results in an interject.

The routine exercise capability allows the DIF the opportunity to functionally verify the operation of critical hardware segments routinely. This constant monitoring provides early error-detection and correction.

VII. CONCLUSION

The development of the DIF was the result of the concerted effort of the components, switching, and transmission organizations in Bell Laboratories and of the components and system organizations in Western Electric. A comparison of the DIF with the DT/SP2, which it replaces, indicates a significant reduction in cost, power, and space requirements, an improvement in reliability, and a reduction in the installation effort. By the end of 1980, it is expected that over 500,000 No. 4 ESS terminations on DIF will have been deployed.

REFERENCES

1. J. F. Boyle et al., "Transmission/Switching Interfaces and Toll Terminal Equipment for No. 4 ESS," B.S.T.J., *56,* No. 7 (September 1977), pp. 1057–97.
2. T. J. Cieslak et al., "Software Organization and Basic Call Handling," B.S.T.J., *56,* No. 7 (September 1977), pp. 1113–38.
3. W. L. Harrod and A. G. Lubowe, "The *BELLPAC*™ Modular Electronic Packaging System," B.S.T.J., *58,* No. 10 (December 1979), pp. 2271–88.
4. R. A. Bruce, P. K. Giloth, and E. H. Siegel, Jr., "No. 4 ESS—Evolution of a Digital Switching System," IEEE Trans. Commun., *COM-27,* No. 7 (July 1979), pp. 1001–11.
5. T. W. Anderson et al., "No. 4 ESS—Mass Announcement Subsystem," B.S.T.J., this issue.