No. 1 ESS Bus System

By J. B. CONNELL, L. W. HUSSEY and R. W. KETCHLEDGE

(Manuscript received January 15, 1964)

Communication between the various units within No. 1 ESS demands a complex network of transmission lines. The requirements are radically different from previous offices because a large amount of information must be transmitted at high speed, in digital form, between a multiplicity of locations with a high degree of reliability.

A description is given of objectives and of organization and operational modes for the network. The problems involved are discussed together with the hardware by which the objectives were successfully realized.

I. INTRODUCTION

The function of the bus system of No. 1 ESS is to provide intramachine data and control communication. The need for large amounts of such communication is a natural outgrowth of the functional unit structure of the over-all system. At the same time the interconnection method must be highly flexible to permit easy growth and to adapt to large variations between different installations.

Existing electromechanical switching systems send most of their intramachine data by dc ground-return signaling. Most relays used in telephony are relatively easy to control at a distance because of their limited frequency response and high signal threshold. Semiconductor logic circuits, on the other hand, have a wide frequency response and can be disturbed by relatively small noise signals.

It became apparent quite early in the development that conventional interconnection techniques would not be satisfactory. The No. 1 ESS bus system described in this paper has proved to be a successful solution of the interconnection problem.

1.1 *Size*

Telephone switching systems tend to be relatively large assemblies of equipment consisting of many separate functional units. Compared to

previous systems, No. 1 ESS is physically smaller and is composed of fewer separate functional units. However, even a small 5000-line No. 1 ESS might typically contain 3 temporary memories, 2 permanent memories, 2 central controls, 2 master scanners, 2 central pulse distributors and a master control center plus 19 switching network, trunk or junctor units — a total of over 30 separate functional units. A very large office would contain hundreds of such units. Similarly, while distances between units may be only a few tens of feet in a very small office, a very large office will require interconnecting leads hundreds of feet long.

To assemble the functional units into a working system requires interconnections which can transmit data and control information between units with the utmost dependability. In most cases a multiplicity of sources and sinks is involved. For example, either central control must be capable of sending orders to hundreds of network controllers.

1.2 *Duplication*

All system units required to provide continuous service are duplicated. This includes the bus system. When one or more functional units are out of service the effective interconnection pattern must be modified appropriately.

This requires that the buses provide suitable interconnections both normally and under trouble conditions. The bus structure must permit assembly of a fully operational system in the presence of any pattern of faults that does not include simultaneous failures of both units of a duplicate pair.

Through the use of duplicated buses and provision of multiple-source, multiple-sink capability, the buses become a means for achieving the necessary duplicate switching.

The Morris, Illinois, trial[1] demonstrated that, to avoid any disturbance of telephone service when a functional unit fails, the duplicate switching must be performed rapidly. Relays used in the Morris trial equipment to switch service to duplicates were found to be too slow. Thus duplicate switching at electronic speeds appeared necessary. To achieve this goal, the bus system provides all interconnection facilities that the system will require under any conditions. At any instant only a limited set of these possibilities will be in use. Each functional unit can be instructed to receive on a particular bus or its duplicate, to send on a particular bus or its duplicate, or in some cases to send on both. Means are also provided to disable the sending circuits so that false pulses due to trouble conditions cannot destroy the usefulness of the bus to the other units it serves in multiple.

1.3 *Speed*

Cycle times for the control actions of No. 1 ESS are 5.5 microseconds. In a typical cycle, a read operation in the program store can occur simultaneously with a read or write operation in a call store, and both of these may be simultaneous with an instruction to a peripheral unit. To keep the time wasted in propagation to a minimum, the bus lengths are kept as short as possible and the buses are word organized, with parallel transmission of bits. Maximum lengths of 125 feet between any central control and any store keep propagation times to a small but not inconsequential fraction of a cycle. For peripheral units, longer bus lengths of up to 450 feet are required in large offices.

Considerations of speed are also involved in the choice of a 0.5-microsecond pulse as the basic bus signal. Shorter pulses would be more difficult to generate, transmit and detect, while longer pulses would cost time. Also, the 0.5-microsecond pulse is short compared to the repetition time of 5.5 microseconds, and dc restoration is therefore not necessary.

1.4 *Environment*

A number of environmental factors complicate the bus problem. These include the use of common storage battery power and the large physical size of the system, which together cause significant dc potential differences to exist between the grounds of various units. This ground potential problem is avoided by making the bus system ac-coupled. This takes care of dc noise, but other noise sources are also important. Relay circuits in No. 1 ESS are carefully protected to minimize noise from this source. However, a No. 1 ESS may be adjacent to an electromechanical central office whose relays might cause interference. Balanced transmission, shielding where necessary, and separation of bus conductors from other leads minimize such noise pickup. In addition, buses generally include synchronizing or enabling signals that reduce noise effects by time discrimination. An enable pulse is used to activate the other bus sensing circuits. Occasional noise pulses occurring in the absence of a pulse on the enable lead will not normally be sensed.

1.5 *Standardization with Flexibility*

A strong effort towards standardization and code minimization has been made in the No. 1 ESS development. Thus the same techniques, circuits and components have been used for all of the high-speed interconnections of the system. However, flexibility of system arrangements has been retained. This is largely due to the multiple-source, multiple-

sink capability of the bus system and to the use of separate program store, call store and peripheral unit buses. This flexibility is particularly important when additions are made to a working office. When a new equipment unit is added, the appropriate buses are easily extended to the new unit. The bus duplication is used to make the extensions without service interruption.

II. INTERFRAME COMMUNICATION SYSTEM ORGANIZATION AND OPERATION

2.1 *General Organization and Growth*

A simplified diagram of the No. 1 ESS System is shown in Fig. 1. The input-output units represent a wide range of physical units: the space-division network controllers, ferrod scanner units, trunk and junctor signal distributor units, message accounting tape units, and teletype-writers. In its simplest form, the central control complex is capable of only one operation with these equipments during any particular input-output cycle. As a result, all such units are designed to work from one master bus system called the "peripheral bus system." This peripheral bus system is expandable to a very large number of input-output units.

Fig. 1 shows duplicate central control units connected to the above mentioned peripheral bus system. It also shows them connected to a group of program stores via a program store bus system and to a group of call stores via a call store bus system.

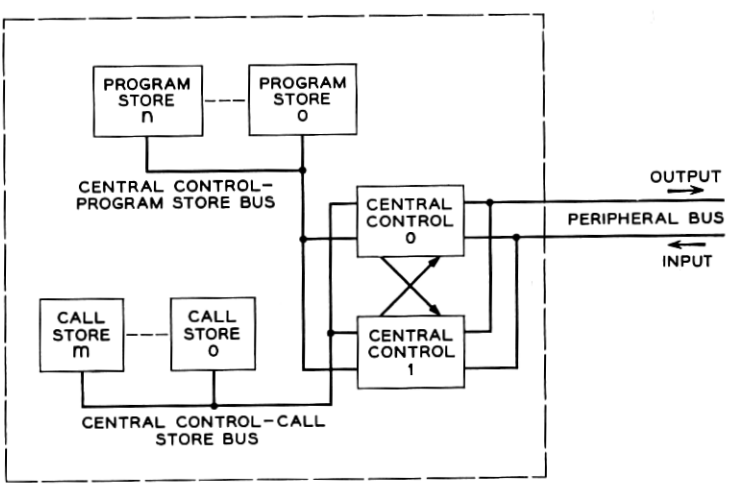

Fig. 1 — Simplified system diagram.

All of the central control operating programs and some rarely changed data (translation data, primarily) are maintained in the program stores. The call store contains data normally modified as the result of processing. By using, with some exceptions, separate memory systems for instructions and for data it is possible to utilize parallel memory operations and instruction overlap operation. This permits more data processing per unit of real time, which in turn increases the call handling capacity of the system.

There are groups of leads interconnecting the central controls shown in Fig. 1. The central controls operate in parallel, and these leads are required for synchronization and for maintenance matching.

2.2 *Bus Control*

Each of the three major bus systems can be thought of as a separate subsystem. The major difference between a bus subsystem and any of the other subsystem units, for example a program store, is that the bus is physically distributed over many equipment frames and its control, although conceptually centralized, is similarly distributed over many units. All bus operations are initiated by the central control. It is the central control that requests information from the program store, reads or writes in the call store, and requests input information or transmits instructions to units on the peripheral bus. Fig. 2 depicts the basic problem for a typical bus system. The central controls are shown on the left of the figure. The buses are shown duplicated, and each consists of two one-way groups of twisted wire pairs (designated "address" and "answer" groups). Each one-way group contains one or more functional groupings, although no such breakdown is shown on this figure. The operation of the bus is such that only one unit (source) can transmit on any bus

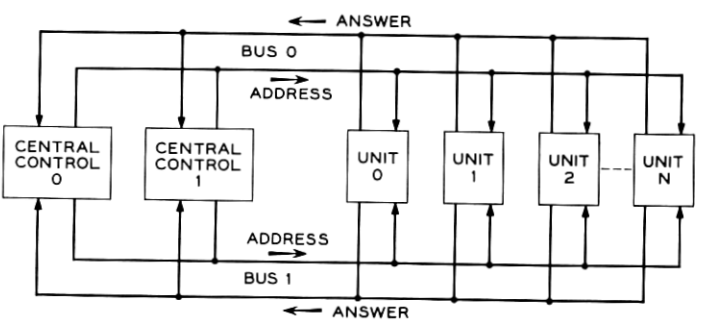

Fig. 2 — Typical bus configuration.

group (e.g., bus 0 address group). However, all the sink units can receive from the same bus. The control problems common to all the No. 1 ESS bus systems are the following:

(1) control of the configuration between central controls and the buses for both transmission and reception,

(2) selection of the unit on the bus for which the transmission is intended, and

(3) control of the configuration between the units on the buses and the buses for both transmission and reception.

The solution to these three problems depends very heavily on the duplication scheme used with the equipment served by the bus system.

2.3 *Program Store Interconnections and Operation*[2]

Due to two major factors, the frequency of use of the memory units and the memory address limitations of the central control, a code select method is used with the memory bus systems. Each transmission from the central control to either the program store or call store bus system is accompanied by a group of code bits (4 bits with the program store bus system and 6 bits with the call store bus system). These code bits are received by all memory units on the bus system. Each store has preset into it the codes for the two information blocks its contains. Only if a store contains the block requested in either its left or right half will it respond to the address bits. The address bits are used to specify the location in the memory block. Normally, two stores will respond to each request.

The addresses for both the program stores and the call stores are generated by the central control. These addresses are formed from a 21-bit memory address field. Table I lists the memory address spectrum assignments. If a program store is being addressed, 20 bits of the memory address spectrum are required. The 21st bit is used to determine which part of the program store word is to be used by the central control. This bit is not sent to the program stores, and it is used only when receiving data from the program store. Of the 20 bits that are transmitted as a program store address, 4 are used as information block code bits and 16 are used to specify an address within the information block.

The program store is a read-only type of memory unit. As a result, there is no data bus for writing into the program store. However, there are other bits of information sent to the program stores. Four bits are used to specify one of five program store operating modes. A synchronizing bit is also sent to the program stores. It is used to control the gating of the information bits into the program store receiving circuits. The

TABLE I — MEMORY ADDRESS SPECTRUM

	P.S. Bus Code Bits		C.S. Bus Code Bits							
20	19	18	17	16	15	14	13	12	11–0	
0	0	0	0	0	0	0	0	0	all 0's	buffer reg operations
0	0	0	0	0	0	0	0	0	all 1's	
0	0	0	0	0	0	0	0	1	all 0's	central control call stores
0	0	0	1	0	1	1	1	1	all 1's	
0	0	0	1	1	0	0	0	0	all 0's	reserved for future allocation
0	0	0	1	1	0	1	1	1	all 1's	
0	0	0	1	1	1	0	0	0	0's–1's	
0	0	0	1	1	1	0	0	1	0's–1's	
0	0	0	1	1	1	0	1	0	0's–1's	
0	0	0	1	1	1	0	1	1	0's–1's	
0	0	0	1	1	1	1	d	d	0's–1's	standby central control
0	0	1	0	0	0	0	0	0	all 0's	right half program store words
0	1	1	1	1	1	1	1	1	all 1's	
1	0	0	0	0	0	0	0	0	all 0's	unused
1	0	0	1	1	1	1	1	1	all 1's	
1	0	1	0	0	0	0	0	0	all 0's	left half program store words
1	1	1	1	1	1	1	1	1	all 1's	

use of synchronizing pulses reduces the time during which the program stores are exposed to noise on the buses.

The central control receives 44 information bits and an all-seems-well signal from the program store. The all-seems-well signal is returned by the program store if certain conditions are satisfied during the execution of the central control request. A synchronizing signal is also sent to the central control together with the program store readout and all-seems-well signal. This is used to reduce the time that the central controls are exposed to noise on the bus system. Fig. 3 depicts the bus system and points out its ability to expand. Also shown in Fig. 3 are two separate answer buses, an east and a west. The need for two separate answer buses arose because of timing considerations and because of the directionality of the cable receivers used in the bus system. The address bus does not need to be separate, since the cable drivers transmit along the bus leads in both directions.

Fig. 4 presents a timing diagram for program store bus operations. This timing diagram shows the relative time between transmissions and receptions. A diagram showing the relative timing between program store bus operation, call store bus operation and peripheral bus operations will be presented later in this article.

Route control flip-flops located in the central controls and program

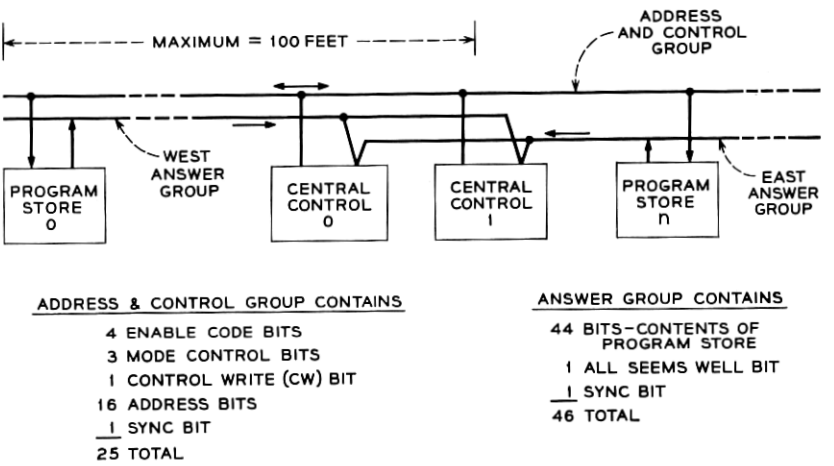

Fig. 3 — Program store bus system.

stores allow a multiplicity of program store bus configurations. Three flip-flops are used in the central controls for normal control. Seven flip-flops within each program store control the configuration between the buses and the store.

2.4 *Call Store Interconnections and Operation*[3]

A central control can address its own call stores, the other central control, or a group of special control flip-flop locations termed "buffer control registers." These are all addressed via the central control index adder output register, receive data from the central control data buffer register, and transmit data to the central control data buffer register. Choice among these actions is dictated by mutually exclusive addresses

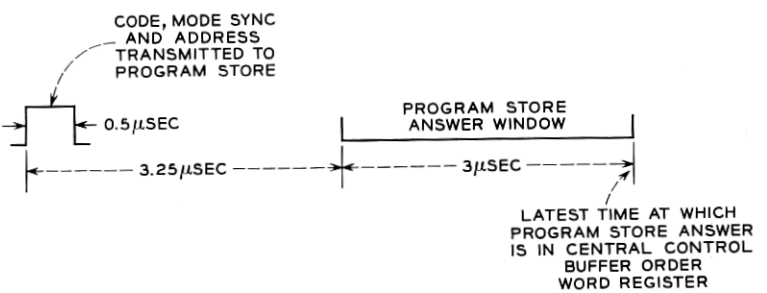

Fig. 4 — Program store system timing.

recognized at the output of the index adder output register in the central control. Table I contains the address assignments for these operations. The buffer control registers are internal to the central control and are not accessed via the call store bus system except when one central control is writing in one of these registers in the other central control. They are mentioned here for completeness.

The call store bus is duplicated. Only one bus is shown in Fig. 5. As shown in this figure, this bus contains an address and control group, an information group for writing in a unit on the call store bus and an answer group for reading from a unit on the call store bus. The address and control group transmits two synchronizing pulses, six enable code bits, three mode-control signals (C, G, H), a 12-bit address, read or write signal and a parity bit which is computed over the address and code. Two synchronizing pulses are required, since the information is sent during two separate time periods. The information group transmits a synchronizing pulse, a 23-bit information word, and a parity signal which is computed over the information, address, and code. This parity bit is stored with the word and is checked later when the word is read out of the call store. Fig. 6 shows a timing diagram for call store write and read operations. The answer bus returns a synchronizing pulse, the stored

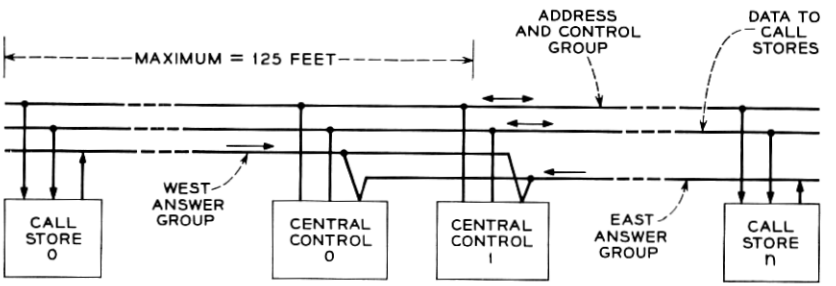

ADDRESS & CONTROL GROUP CONTAINS

6 ENABLE CODE BITS
3 MODE CONTROL SIGNALS
12 ADDRESS BITS
1 READ BIT
1 WRITE BIT
1 PARITY BIT
2 SYNC BITS
26 TOTAL

ANSWER GROUP CONTAINS

24 DATA BITS
1 ALL SEEMS WELL BIT
1 SYNC BIT
26 TOTAL

DATA TO CALL STORE GROUP CONTAINS

24 DATA BITS
1 SYNC BIT
25 TOTAL

Fig. 5 — Call store bus system.

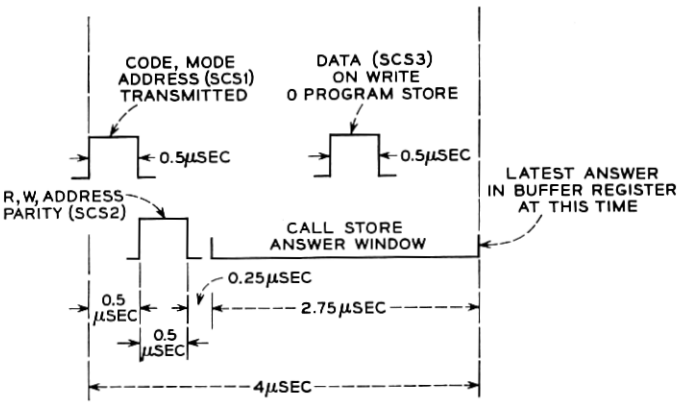

Fig. 6 — Call store system timing.

parity bit, 23 bits of information, and an all-seems-well signal. The all-seems-well signal returns only if certain conditions are met while the call store is executing a central control requested operation. There are east and west answer buses similar to those discussed with the program store answer bus. They are required for the same reasons.

Route control flip-flops located in the central controls and call stores allow a multiplicity of call store bus configurations. Three flip-flops are used in the central controls for normal control. Seven flip-flops within each call store control the configuration between the buses and the store.

2.5 *Peripheral System Interconnections and Operation*

All peripheral units will receive either data or instructions from the duplicated peripheral address bus. Peripheral units such as scanners, signal distributors, and network switch bays have duplicated controllers. Each of the two controllers can be connected to either address bus. The basic bus-to-controller logic for all such units is shown in Fig. 7. The four leads marked E_0, E_1, E_2 and E_3 are the enable leads. They activate the unit and simultaneously select the bus-to-controller path. These four signals are supplied from central pulse distributor units, which are duplicated. Two enable leads are supplied from each of the duplicate central pulse distributors. There are some units on the peripheral address bus which have unduplicated controllers. These controllers will have access to both address buses. For such units, two enable paths, one from each central pulse distributor, will be used to activate the unit and select the bus-to-controller path.

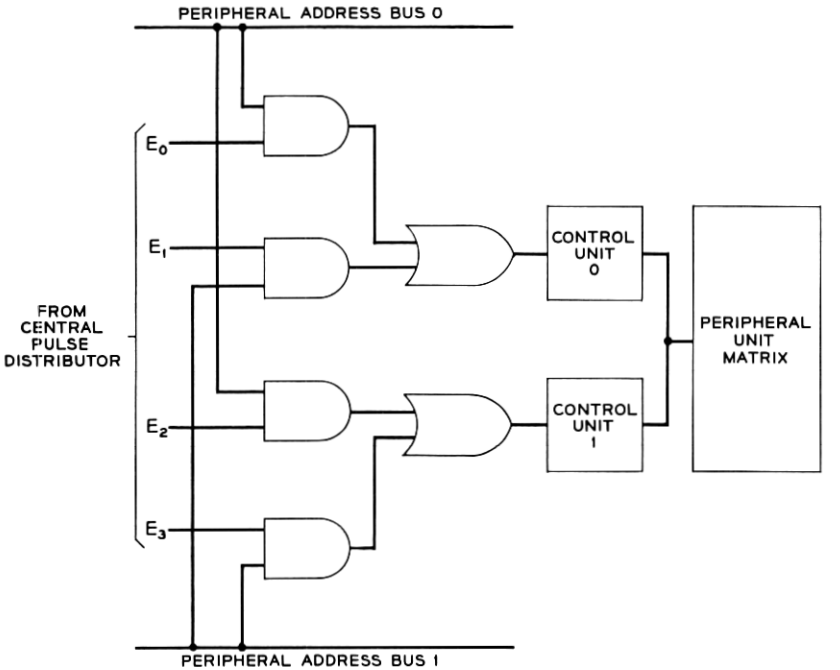

Fig. 7 — Peripheral unit enabling.

In addition to the address bus and the enable facilities for peripheral units, there is also an answer bus for units which transmit information to the central control.

The central pulse distributor outputs are used to perform one of two major functions. The first of these is the peripheral unit enable function and the second is that of providing high-speed discrete control over various flip-flops that may be located throughout the system. The enabling function is performed concurrently with peripheral address operations. As a result, a separate bus system is needed between the central controls and the central pulse distributors. When a non-enable operation is performed using the central pulse distributor, only its bus system is used. Although the central pulse distributors have a separate bus system, it is thought of as being nested into the peripheral bus system and is considered for purposes of operation as well as discussion part of the peripheral bus system.

Fig. 8 is a diagram of the interconnections between the central controls and the peripheral system, including the central pulse distributors. Shown in this figure are:

(1) peripheral unit address bus
(2) peripheral unit answer bus
(3) central pulse distributor address bus
(4) central pulse distributor verify-answer bus
(5) central pulse distributor execute leads
(6) central pulse distributor echo leads
(7) central pulse distributor unipolar outputs (enable leads)
(8) central pulse distributor bipolar outputs.

The central pulse distributor is a matrix of 1024 high-speed, low-level pulsing sources. It can provide two types of outputs; bipolar pulses and unipolar pulses. A bipolar pulse uses both polarities on a single twisted wire pair, whereas a unipolar pulse uses a single polarity on a twisted wire pair. Although the bipolar output is on a single twisted wire pair, it

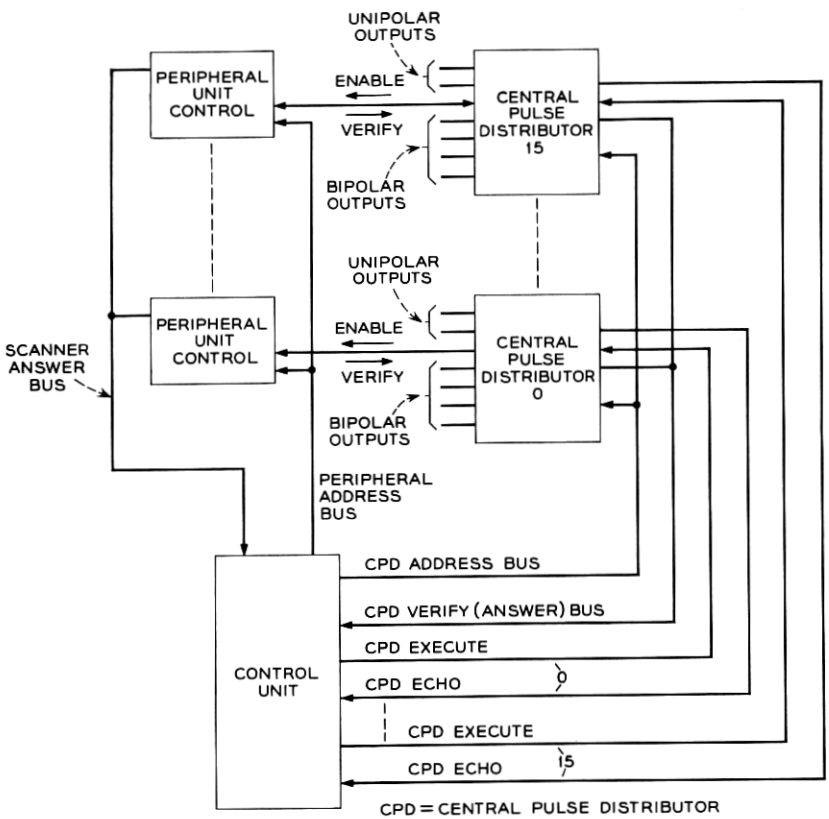

Fig. 8 — Peripheral bus system.

requires two central pulse distributor matrix points to produce the two polarity pulses. One point in the matrix is required for each unipolar output.

The central pulse distributor is effectively a one-out-of-n translator for enabling purposes. As shown in Fig. 8, there can be as many as 16 central pulse distributor units. They are all accessed over a common address bus and they reply over a common answer bus. An execute translator located in the central control is used to select a single central pulse distributor. There are 16 outputs in this translator, one for each central pulse distributor. Upon reception of an execute pulse, the central pulse distributor will transmit an execute-verify signal to the central control over a separate lead. This signal is called the "central pulse distributor echo." Sixteen echo inputs, one from each central pulse distributor, connect to the central control. These signals are received and compared against the output of the execute translator whenever a central pulse distributor is used.

A central pulse distributor has a variable range of combinations of unipolar and bipolar outputs. A maximum of 512 unipolar points can be served by one central pulse distributor. This uses one half of its output matrix. The other half must be used for bipolar pulsing only and can produce 256 such outputs. The unipolar half can be adapted for bipolar use, thus affording up to 512 bipolar outputs per central pulse distributor.

Unipolar pulses are the only type used for peripheral unit enabling. The half of the central pulse distributor matrix that can produce unipolar outputs is equipped with a verification feature. Whenever a peripheral unit is being addressed via the peripheral address bus, an enable address corresponding to that unit must be sent to the central pulse distributor. This address chooses a unipolar output point, and a pulse is transmitted by the central pulse distributor on a single twisted wire pair to a peripheral unit. It is this pulse which enables a peripheral unit to receive information from the peripheral bus. After the information is received from the peripheral bus and checked for proper address, a verify pulse is returned to the central pulse distributor on the same twisted wire pair. The central pulse distributor receives this verify pulse and codes the matrix point on which it was received into the same address code format initially sent to the central pulse distributor from the central control. This is then transmitted to the centrol control on the central pulse distributor verify answer bus. The central control compares this answer to the address it had sent out. If a unipolar output is used for other than peripheral unit enabling and does not return a verify pulse, it is up to the central control to recognize this as the case and ignore the

central pulse distributor verify answer. The verification feature does not exist for the bipolar half of the matrix nor for those normally unipolar half outputs which are connected to provide bipolar outputs.

Binary information is used to select the proper central pulse distributor, address the central pulse distributor and address peripheral units. However, it is not transmitted to the peripheral units or the central pulse distributors in this form. It is pretranslated in the central control.

There are 14 binary bits used to address central pulse distributors. Four of these bits are used to select which central pulse distributor will be used. The ten binary-coded address bits are pretranslated by the central control into two coded groups of one-out-of-eight and one group of one-out-of-sixteen. These require 32 address bus leads. The one-out-of-sixteen group can be thought of as two groups of one-of-eight where one and only one of these groups is used for an operation. It is the division of this group that determines which half of the output matrix is being used (e.g., verification half or not). Thus only three groups of one-out-of eight are required to specify a point in the verification half of the output matrix. The answer bus then consists of 24 leads treated as three groups of one-out-of-eight. Although more address bus leads are required with this pretranslation technique, the saving of translation equipment at the central pulse distributors is sufficient to warrant this operation.

A multiplicity of units share the peripheral address bus system. The major units — scanners, signal distributors and network switch frames — have their information pretranslated at the central control before being placed on the peripheral bus. The information for these units (addresses and instructions) is maintained in the central control in binary form. One register location is used for generating peripheral addresses in the central control; this is the addend K register. A group of translators connect between this register and the peripheral address bus. Selection of the proper translator must be made with every peripheral operation. Fig. 9 shows the information groupings for the units on the peripheral bus. Fig. 10 shows the binary form as placed in the addend K register. Information can also be placed on the peripheral address bus in untranslated binary form. Choice among translators will be discussed below. The size of the peripheral bus in terms of the number of leads is dictated by the translator coding for the network switch bays. As seen from Fig. 9, 36 leads are needed in the peripheral address bus for these units. There are two additional leads in the peripheral address bus; these are the network reset lead and the false cross and ground test lead.

The peripheral answer bus is made up of 17 leads; 16 of these carry answer information, while the 17th is an all-seems-well scanner lead.

The control of the peripheral bus system must be flexible, because of

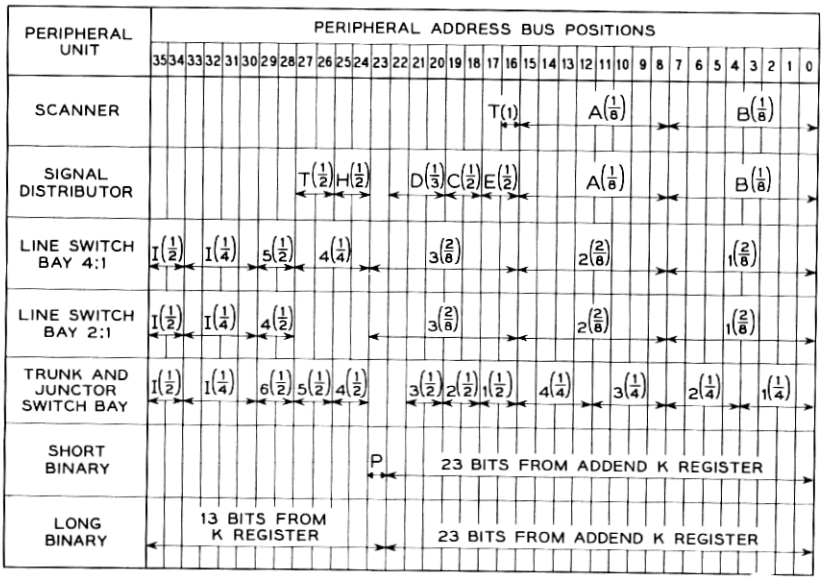

Fig. 9 — Peripheral address bus bit assignments.

the nonhomogeneity of the units connected to it. A given peripheral bus operation takes two machine cycles (11 microseconds). Subsequent peripheral operations could go to entirely diffe.ent types of units: in fact, bipolar (nonperipheral unit) operation could alternate with peripheral unit operation. Thus each peripheral bus cycle can and must be completely independent of any preceding or subsequent operation. The manner in which units are selected, i.e., unipolar enable (see Fig. 7), also determines the bus configuration with respect to the peripheral unit. This is unlike the program and call store bus configuration control. In those systems the bus configuration is independent of the unit selection and the route does not change with the memory unit address. The peripheral system selects the bus configuration while selecting the peripheral unit. The bus configuration between the central controls and the central pulse distributors is selected by one flip-flop located in the central control. This is the CPDB flip-flop, and once set it remains set for many operations until changed via a maintenance configuration control program. At the start of peripheral operations the state of this flip-flop is transmitted to all central pulse distributors, telling them which of the two buses will be used for a subsequent address transmission.

The central control F register usage is the key to the peripheral system control. For each peripheral operation, two binary words must be used

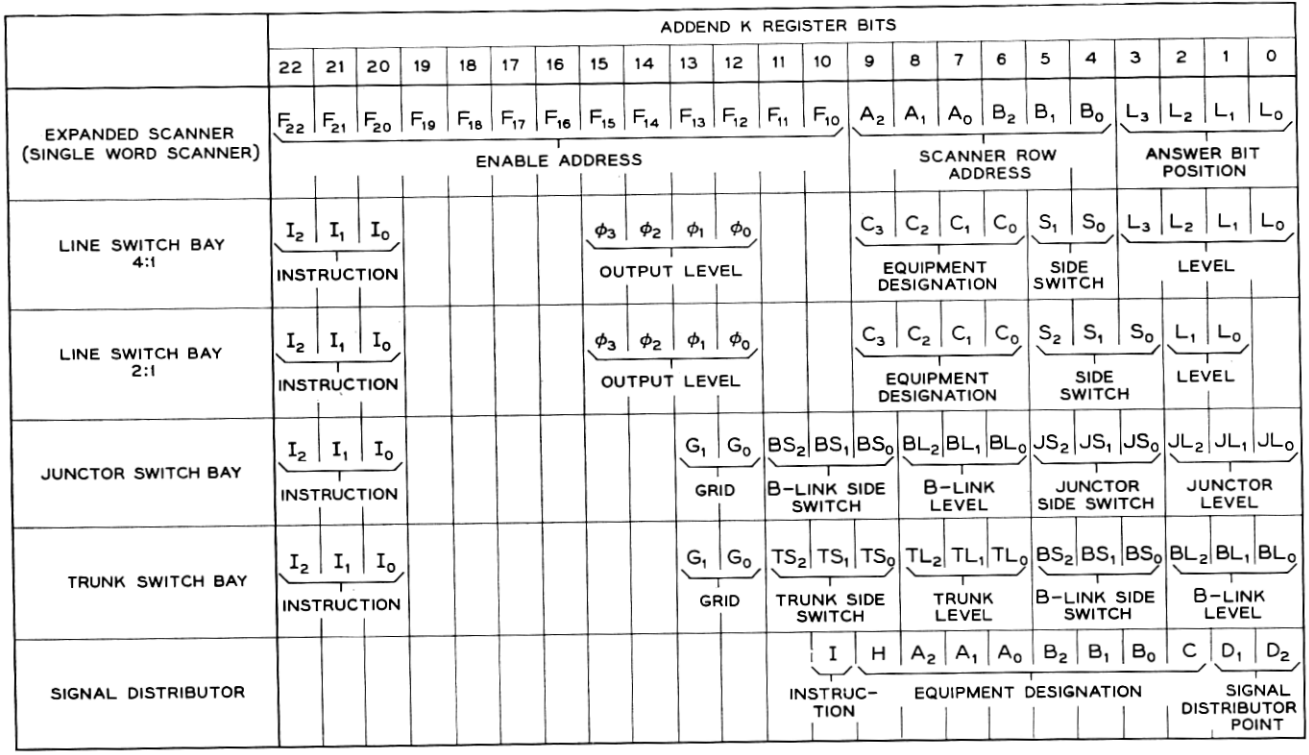

NOTE: BLANK SPACES ARE DON'T CARES

Fig. 10 — Addend K register bit assignment.

by the central control. One is the internal instruction for the peripheral unit and is placed in the central control addend K register; the other word is the peripheral unit enable word and is placed in the central control F register. On nonperipheral unit operations, when only a central pulse distributor output is to be activated, the address of the point is placed in the central control F register and the addend K register is not used.

The F register word contains 23 bits. Bits F_{10} through F_{13} control the execute translator and select the proper central pulse distributor. Bits F_{14} through F_{22} combined with F_9 are the ten binary bits from which the central pulse distributor point address is generated. The use of the F_9 bit is not direct. It becomes part of the address only for non-enable operations. For enabling operations, F_9 is considered equal to zero regardless of its actual value. For non-enable operations, it assumes its actual value. This bit is the one used to select which half of the central pulse distributor matrix is used.

Those central control instructions used for peripheral operations are treated as enable or non-enable instructions. The central control instruction decoder must determine which it is and whether or not a peripheral answer is expected if it is an enable-type instruction. Whenever a peripheral answer is expected, a special flip-flop designated scanner request flip-flop (SCR-FF) is set. When no answer is expected, it is reset. Whenever a non-enable operation is to be performed, another special flip-flop designated the F inhibit flip-flop (FINH-FF) is set. It is also set for a class of single-word scanner instructions (see Fig. 10). It is reset for all other peripheral instructions. The F_9, F_8 and F_7 bits are used to select the peripheral address translator for all enable operations except the single-word scanner instruction. The codes for translator selection are given in Table II. The F_6 bit is used to determine whether or not a verify address is expected from the central pulse distributor. The F_5 bit is used

TABLE II — TRANSLATOR SELECTION CODING

F_9	F_8	F_7	Translator
0	0	0	short binary to peripheral bus
0	0	1	long binary to peripheral bus
0	1	0	line switch bay, 4:1 concentration
0	1	1	line switch bay, 2:1 concentration
1	0	0	junctor or trunk switch bay
1	0	1	long binary to CPD address bus
1	1	0	signal distributor
1	1	1	scanner

to determine whether or not to expect an all-seems-well scanner response. The F_4 through F_0 bits are reserved for maintenance use and exert no control over peripheral operations. Upon detecting a peripheral operation, the central control instruction decoder either sets or resets the SCR and FINH flip-flops and starts a peripheral control sequence circuit. This sequence circuit will control all peripheral operations. This sequence circuit uses the SCR and FINH flip-flops as well as many of the F register bits to control peripheral operations. The use of the F register as functions of the SCR and FINH flip-flops is as follows:

(1) FINH = 0 and SCR = 0 or 1, use the F_9, F_8 and F_7 bits to select the peripheral translator,

(2) FINH = 1 and SCR = 1, use the scanner translator for the peripheral address bus, ignore the condition of F_9, F_8 and F_7 bits with regard to selection of the translator. For this particular case, as can be seen from Fig. 10, these bits are part of the scanner address.

(3) FINH = 1 and SCR = 0, no use is made of the peripheral address bus, and the F_9 bit is used as part of the central pulse distributor address selection.

The bus-to-controller path for a peripheral unit is selected by the enable pulse that it receives. A typical two-controller unit has four possible enable pulse inputs, two from each central pulse distributor of a duplicate pair (see Fig. 7). Bit position F_{10} of the central control F register is the least significant of the four central pulse distributor select bits. It is used to select between a duplicate pair of central pulse distributors. The F_{14} bit of the same register is used to select between halves of the enable portion of the central pulse distributor output matrix. Bits F_{10} and F_{14} control the selection of the one-out-of-four unipolar enabling distributor output points, two of which are located in each central pulse distributor of a duplicate pair. Table III shows the route selection for all such peripheral units. From this table it is clear that F_{14} is also used to select the proper address bus for transmission.

The F register word is called the "enable address," since it chooses the address bus and indirectly the proper controller. This word can differ

TABLE III — PERIPHERAL UNIT ROUTE SELECTION

F_{14}	F_{10}	CPD of a Pair	Peripheral Unit Controller	Address Bus
0	0	A	A	A
0	1	B	B	A
1	0	A	B	B
1	1	B	A	B

from unit to unit. For this reason, the normal procedure is to "look up" the enable address in a call store memory table. This table is termed the "enable" table.

When performing non-enable operations, the F_{14} bit is still used to control the selection between halves of whichever portion of the central pulse distributor matrix the F_9 bit selects. As such, the F_{14} bit is used to determine whether or not a bipolar output pulse is the set or reset polarity. If a pair of unipolar outputs is used for such control, the F_{14} bit is used to choose between them.

Although the major peripheral bus configuration control takes place via the same mechanism as the unit selection, there are several operating modes which affect this bus configuration. Mode-control flip-flops are used to allow additional operating modes. There are two mode flip-flops in each central control for central pulse distributor address bus control (see Table IV) and two for peripheral address bus control (see Table V).

In order to properly diagnose certain classes of faults, it is necessary to provide independent operation of central controls. This independent mode is referred to as "off-line" operation. The two central controls can work with independent memory units. However, many of the peripheral units are common matrix units, and therefore two nonsynchronized machines cannot be given simultaneous control of the peripheral system. It is still necessary, however, to allow off-line operation of the peripheral system. This is implemented using two flip-flops. These are located in each central control. The central controls react to these flip-flops as shown in Table VI.

Fig. 11 is a timing diagram for peripheral operation. All transmissions in the program store and call store system were accompanied by sync pulses. These are used to reduce the time the receivers on the buses are

TABLE IV — CENTRAL PULSE DISTRIBUTOR MODE CONTROL

CPD Bus Mode Control Flip-Flops		Mode
CDMA	CDMB	
0	0	normal[1]
1	0	mode A[2]
0	1	mode B[3]
1	1	not used

[1] Normal mode: only active control unit transmits over bus designated by CPDB flip-flop.
[2] Mode A: only active unit transmits over both buses.
[3] Mode B: active unit transmits on bus designated by CPDB flip-flop — standby unit transmits on other bus.

TABLE V — PERIPHERAL BUS MODE CONTROL

Mode Control Flip-Flops		Mode
PBMA	PBMB	
0	0	normal[1]
1	0	mode A[2]
0	1	mode B[3]
1	1	not used

[1] Normal mode: only active central control transmits over bus designated by bit 14 of F register.
[2] Mode A: active unit transmits on both buses. Standby does not transmit on either bus.
[3] Mode B: active unit transmits on bus designated by bit 14 of F register; standby unit transmits on other bus.

open and susceptible to noise pulses. The peripheral bus system uses no sync pulses. Some noise protection is afforded, however, through the use of other techniques. When a peripheral unit is enabled by a central pulse distributor unipolar pulse, the mechanism that recognizes the enabling is also used to open a "window." The address for this unit must arrive on the peripheral address bus during the time this "window" is open. The window interval is 2.5 microseconds maximum. Thus, the address and enable signals must arrive at the peripheral unit within a narrow band of time coincidence. Two factors control this coincidence. One is the time of transmission from the central control; this can be strictly controlled. Another is the cabling of the address bus and the enable leads throughout the office. To control the differential between the address and enable route in all sizes of offices, the central pulse distributor frame is used as a distribution center for the peripheral address bus. In this way the distance, and therefore the delay, between the point of transmission, the central control, through the central pulse distributor frame to the peripheral unit is the same for both the address and enable signals. Fig. 12 shows

TABLE VI — OFF-LINE OPERATION

Off-Line Control Flip-Flops		Active Unit	Standby Unit
OL1	OL2		
0	0	E & A	nothing
0	1	E	A
1	0	A	E
1	1	nothing	E & A

E = Enable operation.
A = Address operation.

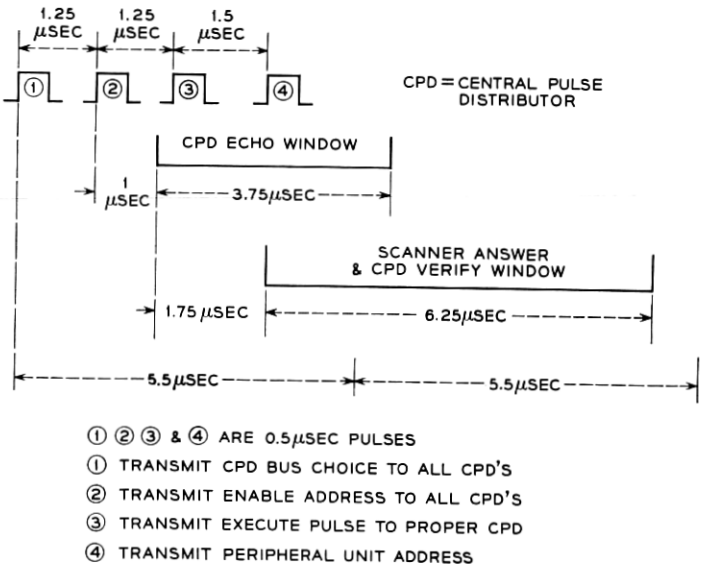

Fig. 11 — Peripheral system timing.

the arrangement in a central pulse distributor. There is provision for four branch address buses in each central pulse distributor. If a second central pulse distributor pair is required for enabling purposes, then the branch address bus used for the units enabled by the additional central pulse distributor pair must emanate from the new central pulse distributor frame.

The peripheral answer bus also passes through the central pulse distributor frames. Within these frames, the answer buses are "fanned in" from a maximum of four to one main answer bus. The main answer bus connects to the central controls. The branch address and branch answer buses give the peripheral system greater flexibility than was required with the program store and call store systems. This flexibility was not needed in these cases since there was a limited number of units on each bus system. With the peripheral bus system, the number of physical frames can run as high as 400 or 500.

With both the address and answer buses, the 0 bus is both fanned in or fanned out at the 0 central pulse distributor of a pair and the 1 bus at the 1 central pulse distributor of the pair.

In the central pulse distributor, there are many output points associated with functions other than the enabling of peripheral units. This was mentioned earlier. Some of these output points control critical

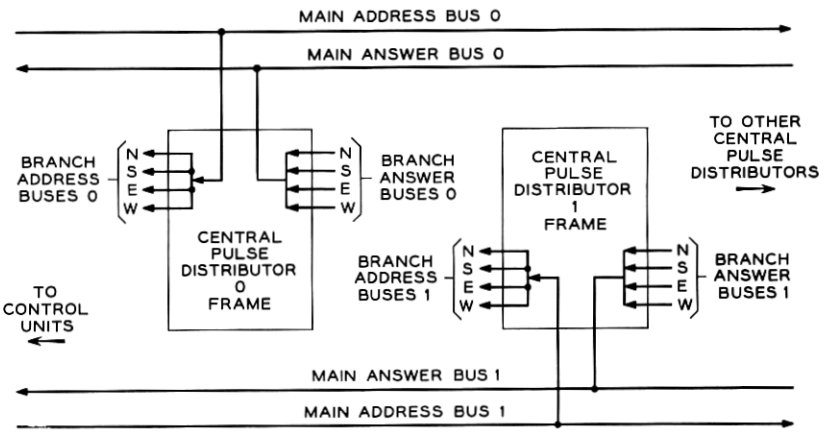

Fig. 12 — Address bus fan-out and answer bus fan-in.

system flip-flops — for example, which central control is active and which is standby — as well as the call store and program store route flip-flops. An unwanted change of these flip-flops could be caused by random noise which may appear on the wires connecting them to the central pulse distributor. A means of protecting such flip-flops has been incorporated into the system. This consists of a common synchronizing signal sent to all such flip-flops simultaneously with the central pulse distributor output pulse. Only the active central control can transmit this pulse. This common signal is designated WRMI ("we really mean it"). It is transmitted from the central control over two twisted wire pairs for duplication and synchronizes in the central pulse distributor frames with the central pulse distributor execute signal. It is fanned out much the same as the address bus.

Fig. 13 shows the relative occurrence of operations on the three bus systems. All call store and peripheral bus operations are controlled by program instructions. The program store bus system is controlled by program instructions only for transfer instructions and data readings.

Because of the necessity for speed and the use of each bus system for round-trip communications, the propagation time, and hence the length, of the buses becomes a critical factor. In the No. 1 ESS system the program store bus is limited to 100 feet and the call store bus to 125 feet from the farthest central control to the farthest memory unit. It was mentioned earlier that two bus cables are provided in both the program and call store bus systems. This is accomplished by combining two unidirectional answer buses at the central controls. The address bus is bidirectional.

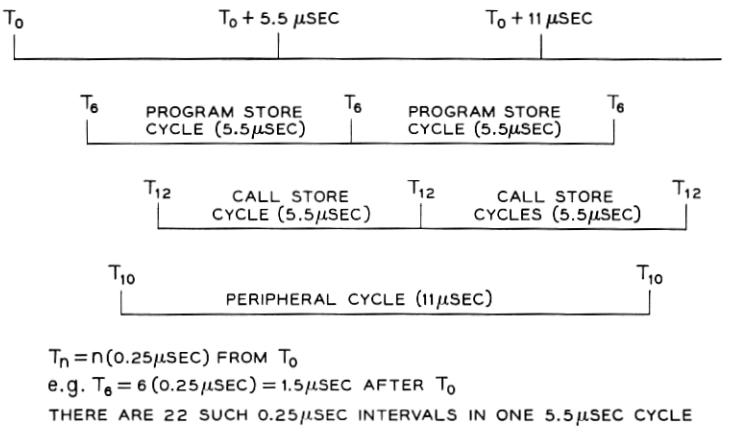

Fig. 13 — Basic communication timing.

The peripheral bus system is used to communicate with many more units than either of the memory bus systems. As a result, longer bus lengths must be used. In addition, two cycles, or 11 microseconds, are used for peripheral bus operations. The distance from the farthest central control through the bus fan-out points to the farthest peripheral unit can be no more than 450 feet. An equally critical factor in the peripheral system is the differential distance between the address bus and the enabling path for the peripheral unit. This differential is minimized by controlling the cable routes.

III. BUS CIRCUITRY

3.1 *The Peripheral Address Bus*

3.1.1 *Typical Pair*

Consider a typical pair in the peripheral address system, starting at one of the central controls. The twisted pair is in a standard switchboard cable and is made of 26-gauge wire with approximately one twist every 2.5 inches.

In a small office one end will go up from a central control to the cable rack, into a special shielded compartment, and over the row of frames to an end guard at the end of the row. Within this end guard are banks of resistors. The address pair terminates in one of these resistors.

The other end goes to the other central control, also by way of the cable rack. After going into the second central control and out again, it

goes, via the cable rack (as always), to a central pulse distributor. Here it has an option. In a large office it would go in and out of a CPD and might go on to do the same at several others. In the small office it is simply connected to a pair of terminals of the CPD, but does not actually go in (see Section 3.7), and then continues in turn to all the peripheral unit frames which contain scanners and network controllers. Beyond the last controller or scanner the other end terminates in a resistor on a nearby end guard.

Before considering the associated hardware it may be noted why a twisted pair was chosen. From the previous discussion, it is evident that this is a transmission line with primary considerations of transmission loss and speed, immunity to noise, reflections and crosstalk, and low cost. One obvious possibility, because of its excellent transmission properties and shielding, is coaxial cabling. However, the cost of coaxial cable is high* and the costs of terminating coaxial wiring are worse. Fortunately, twisted pair, carefully balanced and properly terminated in 100 ohms, has characteristics which, for our purposes, are satisfactory. At the frequencies of interest, the effective delay is 2 μsec/1000 feet. The loss varies with frequency (6 db/1000 feet at 1 mc), producing the distortion illustrated in Fig. 14. As will be noted, tolerable delay limits the length to less than 450 feet where the loss and distortion are relatively small. As will be discussed later, reflections, noise, and crosstalk can be held within satisfactory limits.

3.2 *Grounding Inductance*

Precise balancing to ground is essential to the good noise and crosstalk characteristics required.

The balanced grounding cannot be done satisfactorily by means of the terminating resistors. This would require an expensive pair of matched resistors at the terminals. Instead, a 0.25 mh inductor is used, connected across the pair and grounded at the midpoint. A much better balance is obtained than is practical with resistors, and the ground can be located at the most desirable point near the middle of the longest buses.

The inductor is built into terminal blocks where the pairs are brought in or out of a particular frame.[4]

3.3 *Cable Receivers*

Since this pair transmits a signal from a central control to a large number of receivers in the frames which contain network controllers,

* Even in a relatively small office, the 38 pairs in the address bus alone may add up to about 3 miles of pairs and some 3000 terminations, so interconnections are an appreciable part of the cost of an office.

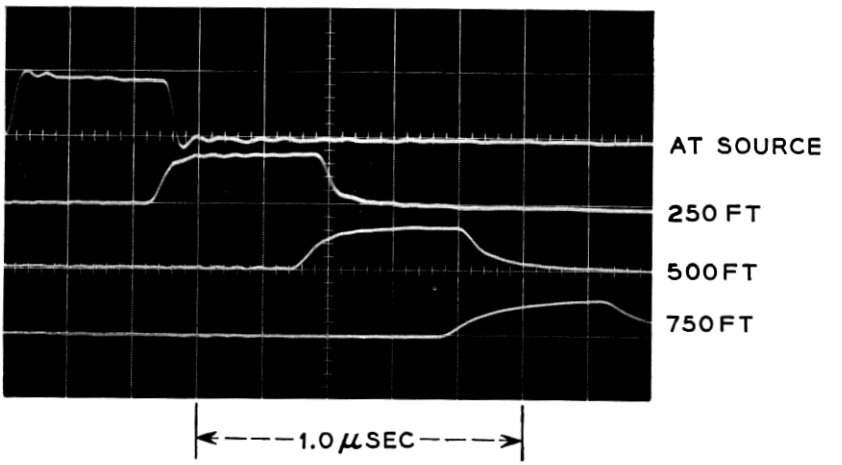

Fig. 14 — Pulse distortion by transmission.

scanners, and signal distributors, it is essential that the individual receivers should not seriously degrade the transmission characteristics of the pair. In addition, a failure in any one receiver should not interrupt transmission to the other locations.

To meet these requirements, a small current transformer is used for pick-off. The primary winding is essentially the bus pair passing through a toroidal core in such a manner as to constitute a balanced, two-turn primary winding. The secondary is a 50-turn winding.

As shown in Fig. 15, the secondary winding goes to a simple grounded-emitter transistor amplifier.

This receiver picks off not more than one per cent of the power being transmitted in the pair and amplifies it to a useful level for use in the frame. This circuit adds a small loss in the balanced pair when operating properly, and nothing which can happen in the output side of one or two pick-offs will seriously damage transmission to other frames. Under normal conditions each pick-off inserts a balanced impedance, less than one ohm, into the line. If the secondary of one transformer is either shorted or opened, the effect on transmissions is a small change in the inserted impedance. In the worst case (open secondary) the inserted impedance in the line is 3 μh. This is harmless if only one or two are open circuited.

3.4 *Cable Drivers*

The standard pulse on the pair is produced by a cable driver. As may be noted in Fig. 16, this is a two-input logic circuit.[5] It is designed to

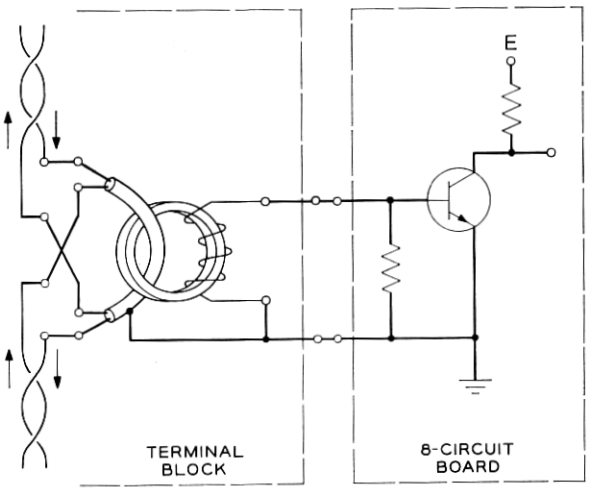

Fig. 15 — Pick-off transformer and cable receiver amplifier.

operate at a higher level than the normal logic circuitry. The two inputs are driven in coincidence by a pulse from standard logic circuits and a timing, or gate, pulse. The output transformer, designed to work into 50 ohms, is connected across the transmission line, driving the bus in both directions.

The requirement that two or more drivers be connected to the same pair and drive at a level adequate for operation of many receivers made the shunt connection the only practical one. The transformer design was the most exacting in this circuit. It was necessary that a physically small

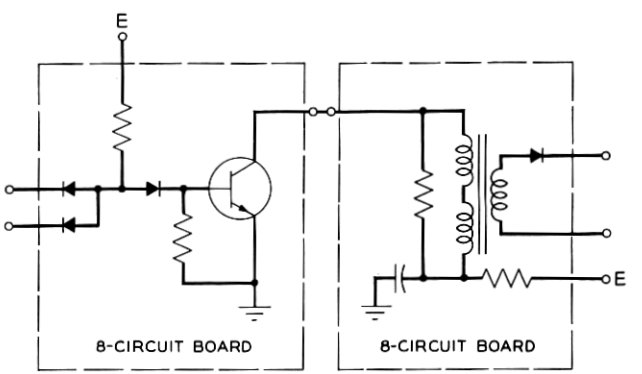

Fig. 16 — Cable driver amplifier and output transformer.

unit (mounted on a standard printed wiring board) meet the requirements: drive a 200-ma, 0.5-μsec pulse into a 50-ohm load on a 10 per cent duty cycle. A 2:1 turn ratio is used, with 3 equal windings, to limit the transistor current to 100 ma. The inductance per winding is 1.35 mh. There were two serious problems. The impedance looking back into the transformer was not large enough to prevent deterioration of the transmission when several circuits were across the same line. This was solved by means of a series isolating diode. A second problem was a sharp reverse spike, caused by leakage inductance, at the end of an output pulse. This spike temporarily increased the transistor collector voltage above breakdown and, in extreme cases, produced noise in other channels where amplifiers were on the same board. These defects could not be cured by decreasing the magnitude of the primary damping resistor without increasing the transistor load and lengthening the transformer duty cycle. The problem was solved by trifilar winding. The result was a slightly increased capacitance (which was desirable here), a leakage inductance cut to about one-third its previous value, and a unit which could be much more readily manufactured to specifications.

3.5 *Noise, Distortion, and Crosstalk*

Before going to more specialized circuitry which is associated with the bus system, the situation on distortion, noise, and crosstalk may be summarized.

As noted previously, delay requirements limit cables to less than 450 feet long. Over that distance the transmission distortion is small. There is one difficulty. While the amplifiers regenerate and transmit a good pulse, they also tend to lengthen it. This becomes serious in a large office where more amplification is used. It will be discussed later.

Putting the buses in a separate section of the cable rack, shielded from the rest of the world, is, in a sense, extra insurance against noise as far as the balanced line itself is concerned. Extensive tests have been made simulating noise (such as that from unprotected relays) which might be encountered. No trouble was detected from the transmission line itself. Longitudinal noise did, however, leak into the cable receiver. The cable receiver transformer transforms from balanced to grounded transmission. At high frequencies the parasitic capacitance ruins the balance, and so a high-level, sharp spike of longitudinal noise could be transmitted through the cable receiver. This problem was solved by shielding the primary winding as shown on Fig. 15.

The greatest remaining source of noise and crosstalk is within terminal

circuitry rather than the bus itself. The circuitry is closely packed, eight circuits to a board on closely packed boards. All circuits under such conditions are prone to noise and crosstalk because of coupling and common ground connections. The amplifier circuits, because of the large currents involved, are particularly so. The solution required careful design (and redesign!) to keep common ground leads short and isolate noisy output leads from sensitive input ones.

3.6 *Grouping of Drivers and Receivers*

Closely analogous to noise problems are those concerned with multiple drive points and multiple reception. As described, an address pair has two or more cable drivers, grouped in the control area, any one of which may drive a multiplicity of receivers in the peripheral area. In a large office there may be up to 50 receivers scattered along each half of the pair with the drivers grouped together in the middle.

There are also answer buses which transmit from a multiplicity of peripheral points to a few receivers in the control area. Here a one-sided bus must be used. Drivers cannot be put on both sides of a receiver. The necessary polarities are such that a pulse from a driver on one side passes through the isolating diode of a driver on the other side in the forward direction. It can temporarily break down the corresponding transistor, or turn it on, producing noise and distortion on the pair. In the case where the address bus goes two directions, two answer buses must be used, going to separate receivers.

3.7 *Central Pulse Distributor*

There are two distinct parts to the central pulse distributor. One[6] decodes an enable address and sends an enable signal on private pairs to other frames. The other part is the one to which the peripheral address pair goes. This section contains a group of amplifiers which regenerate the address pulse and can transmit it simultaneously to a maximum of four outputs. In a small office this amplifier is not needed and this section is wired but not equipped. The address bus is brought only to the CPD terminals so that it will be available if the office grows.

If the expansion in the number of receivers or length of bus exceeds the capacity of one bus, then the CPD is equipped. The pulse is picked off through a standard cable receiver which drives a bus fan-out circuit (see Fig. 17). This circuit simply transformer-couples the input of two transistor amplifiers in parallel to the output of a receiver amplifier. These amplifiers, through standard transformer circuitry, drive two pairs.

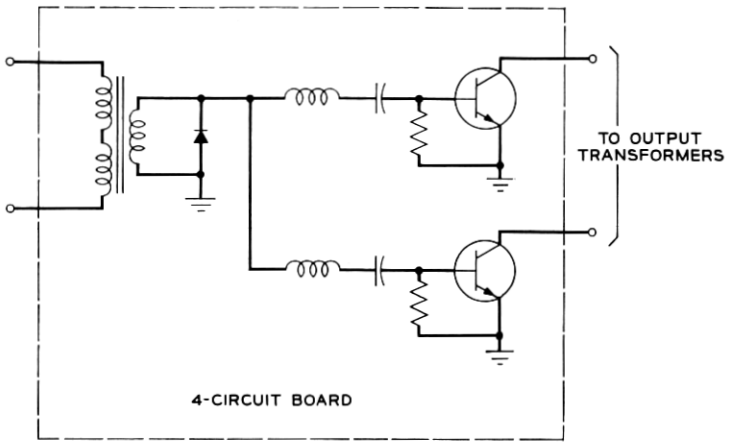

Fig. 17 — Bus fan-out circuit.

Since each end of each bus can be used, this gives fan-out from one bus to four in parallel.

The addition of this amplification results in a difficulty previously mentioned. Each time the pulse is regenerated, the output lengthens a little. At this point there is danger that the accumulated stretching will give a pulse longer than the 0.75-μsec maximum permitted in the peripheral area. To avoid this, the pulse width as well as amplitude is regenerated. *LC* circuits in series with the transistor inputs convert the input pulses to a half sine wave with 0.5-μsec duration. The amplifier clips this, producing a new 0.5-μsec pulse.

One further note may be made on bus length. The delay which may be tolerated is measured from the first cable driver in the control area to the farthest cable receiver. The 450-foot maximum length must be measured in the same way from the control center, not the CPD. Since the corresponding enable pulse must coincide with the address at the receiver, the buses involved must be of equal length, within 50 feet. To insure this, they follow the same path, as far as possible.

IV. OTHER BUSES

4.1 *Answer*

The answer bus sends signals between the same units as does the address bus, but in the reverse direction, so it has many cable drivers scattered along it and a few receivers in the control area. As noted, the

characteristics of the cable drivers make this a one-sided bus — drivers cannot be located on both sides of the receivers. It uses exactly the same circuitry as the address bus, with one exception. In the large office the answer bus must fan in four-to-one. This is done using a combination (Fig. 18) of receivers, logic circuits and drivers in an OR configuration.

4.2 *Enable Address*

The enable address is decoded in its section of the CPD, and the terminating circuitry is somewhat different. The input from the enable address bus to the CPD's is picked off in the manner already described, but the outputs are on a one-to-one basis, each enable pair going to one receiver. These pairs are transformer-coupled in special circuitry[6] within the particular frames. As discussed in the reference, there are several types of these pairs and they may transmit the standard pulses or a bipolar pulse. The enable verify is an exception to the rule that buses are

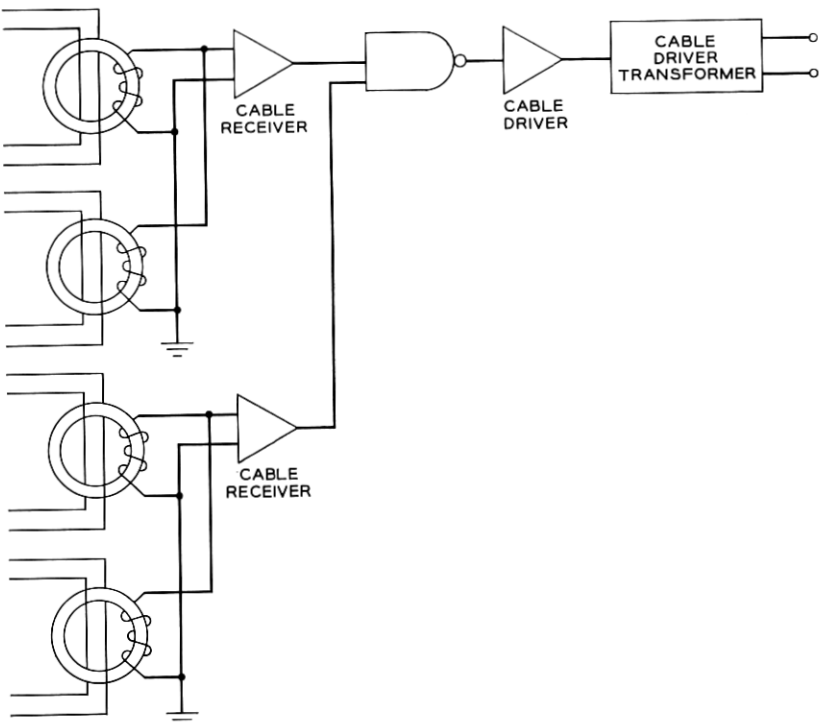

Fig. 18 — Bus fan-in circuit.

all one-way. When an enable pulse is transmitted, a verify is returned (after a brief delay) to the CPD over the same pair, using special circuitry.[6] From the CPD the verify pulse gets to the control center over an enable verify bus in the standard manner.

4.3 *We Really Mean It and Set Manual Pulse Bus*

There are many private pairs, similar to the enable, which are used to enable or disable various sections of the office — generally by operating flip-flops. These are operated by pulses generated in the enable section of the CPD's and the coincident gate (WRMI), which is a simultaneous pulse produced synchronously from a special fan-out circuit. Three of the four outputs are always used even in a small office:

WRMIA goes to master control center, master scanner and central control

WRMIC goes to call stores

WRMID goes to program stores.

All use standard drivers and receivers.

A special fan-out circuit is used because the WRMI pulse must be stretched when it arrives at the CPD before being gated out.

The pulse is received by means of a cable receiver pick-off and amplifier (Fig. 19). This drives a dynamic register[6] which stretches the pulse to 2 μsec. During this interval the four WRMI output pairs may be pulsed coincidentally with the pulse from the enable section of the CPD. This is done by standard cable driver circuitry. The additional cable driver shown on Fig. 19 is needed to correct the polarity.

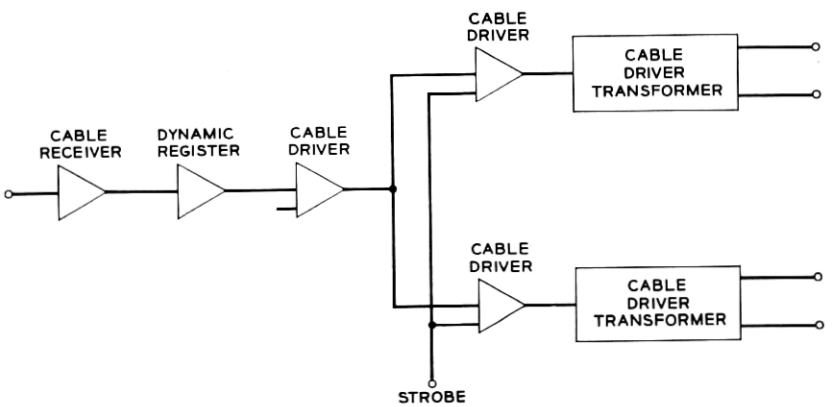

Fig. 19 — WRMI fan-out circuit.

The SMPB pair is discussed with WRMI only because it also is a lone pair — a gate or execute driven from the master control center and having receivers at central controls and program stores.

4.4 *Stores*

The individual enable and status connections from the CPD's to the stores have been mentioned. The other major interconnections in this area are the call store or program store address and answer buses. For timing reasons these must be kept short — not more than 100 feet in length or more than 24 receivers per bus. There is the same sort of limitation as on the peripheral address and answer. Drivers in the control area can operate both ends of the address bus, but the answer bus must be one-way. Drivers at the stores must all be on the same side of the receivers in the control area. Standard cable drivers are used for these buses. Standard receivers are used at the central control. At the stores the cable receiver current transformers are standard, but the amplifiers are part of a more complex circuit.

4.5 *Private Buses in the Control Area*

Most of the pairs between the central controls for maintenance and diagnostic purposes use standard drivers and receivers. There are also timing (clock pulse) pairs which have carefully controlled transmission characteristics. These connect directly to special circuitry within the central controls.

V. DC CONNECTIONS

There are three types of these:

5.1 *Master Scanner*

This collection of connecting pairs is used to observe the status of test points scattered throughout the office. The scanner itself is not very sensitive to noise but, because the leads are long and are exposed to many sources of noise, there is a danger that they might transmit noise to another sensitive circuit. For this reason, twisted pairs are used and the leads are carried in their own shielded section in the cable rack. They are of course isolated from the frame ground, the ground return being brought back to the power ground for that circuit.

5.2 *Trunk Circuitry*

The interconnections for the trunk circuit relays are also twisted pairs grounded in the same manner as the master scanner pairs. They are also, as far as practical, balanced to ac to suppress relay noise.

5.3 *Maintenance and Control Circuitry*

There are a number of relays used for maintenance and control purposes. These are operated by dc connections similar to those of the trunk circuitry.

VI. OPERATING EXPERIENCE

As previously noted, extensive laboratory tests were made, simulating the worst predictable conditions. Troubles were detected (such as the need for shields on the pick-off transformers) and eliminated.

The central control and associated buses were operated in and close to an operating, electromechanical central office (even with protection taken off adjacent relays) without failure.

The most extensive and significant experience is with the actual offices — with the Holmdel, N. J., systems laboratory, which is a small office with a complete interconnecting bus system, and with the first commercial office, installed and now under test at Succasunna, N. J.

Testing of an assembled office was begun in Holmdel in April, 1963. It has progressed to the state that calls have been set up through it. Testing of the Succasunna office started in August, 1963, and is continuing. Most of this testing uses the interconnecting buses, so this part of the office has been thoroughly tested.

These tests prove that this is a working system which should operate satisfactorily in commercial service.

VII. SUMMARY

The complex network of transmission lines described here is organized to meet the objective — the rapid exchange of a large amount of information (in digital form) between many functional units. The organization and operational modes result in flexible, dependable and efficient use of hardware. Simple, reliable hardware is used with a high degree of standardization, resulting in over-all economy and reliability.

REFERENCES

1. Keister, W., Ketchledge, R. W., and Lovell, C. A., Proc. I.E.E., **107**, Part B, Supplement, Nov. 20, 1960.

2. Ault, C. F., Gallaher, L. E., Greenwood, T. S., and Koehler, D. C., No. 1 ESS Program Store, B.S.T.J., this issue, p. 2097.
3. Genke, R. M., Harding, P. A., and Staehler, R. E., No. 1 ESS Call Store — A 0.2-Megabit Ferrite Sheet Memory, B.S.T.J., this issue, p. 2147.
4. Chevalier, J. G., and Eisenhart, R. K., No. 1 ESS Circuit Packs and Connectors, B.S.T.J., this issue, Part 2.
5. Cagle, W. B., Menne, R. S., Skinner, R. S., Staehler, R. E., and Underwood, M. D., No. 1 ESS Logic Circuits and Their Application to the Central Control, B.S.T.J., this issue, p. 2055.
6. Freimanis, L., Guercio, A. M., and May, H. F., No. 1 ESS Scanner, Signal Distributor, and Central Pulse Distributor, B.S.T.J., this issue, Part 2.