THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 56, No. 7, September 1977
Printed in U.S.A.

No. 4 ESS:

Software Organization and Basic Call Handling

By T. J. CIESLAK, L. M. CROXALL, J. B. ROBERTS, M. W. SAAD, and J. M. SCANLON

(Manuscript received July 12, 1976)

The No. 4 ESS has the majority of its control, administrative, and maintenance functions implemented by software. This software is organized into an operating system that provides scheduling and interrupt handling, operational programs that perform the required call-handling and administrative functions, and maintenance programs that provide diagnostic capability for trunks and switching hardware. This paper provides a general description of the design concepts and organization of the No. 4 ESS software, dealing in some detail with the operating system. Functional capabilities, design constraints, and organization of the call-handling software programs are also covered.

I. INTRODUCTION

The No. 4 Electronic Switching System (No. 4 ESS), a stored program switching system, has most of its required control, administrative, and maintenance functions provided via program. The initial generic program for No. 4 ESS is the largest initial program developed to date for Bell System Electronic Switching Systems. It consists of approximately 400,000 instructions and 400,000 words of diagnostic tests requiring in excess of 1.4 million 26-bit words of system storage.

In order to meet the overall design objectives for No. 4 ESS, four major objectives were set for the initial program:

(*i*) Efficiency in real time.
(*ii*) Simple man-machine interfaces.
(*iii*) Defensive design.
(*iv*) Ease of modification.

Meeting these objectives required a unified software/hardware impact on the system architecture and the introduction of a number of improved software design and implementation concepts.[1]

The generic program concept,[2] used in earlier ESSs, will be followed in the introduction of additional features to No. 4 ESS. Each new generic will contain feature additions and will be built on the previous generic's program base.

II. SOFTWARE ORGANIZATION

The No. 4 ESS software consists of three basic categories of programs. They are:

(*i*) Operating system programs, which provide task scheduling, interrupt handling, and man-machine input/output interfaces.
(*ii*) Operational programs, which provide call-handling and administrative capabilities.
(*iii*) Maintenance programs which provide trunk maintenance and frame diagnostics.

Components of both the operating and maintenance categories which are common to the other application of the 1A Processor (i.e., No. 1A ESS[3]) are designed as a separate package called 1A Processor common software.[4] These programs perform the functions of initialization and system restart, fault recovery, and diagnostics for 1A Processor hardware, and input/output for the various I/O devices. This common software represents approximately 150,000 instructions and 215,000 diagnostic tests.

All programs are executed from core (program store or in some cases call store), although their resident storage medium may be tape, disk, or core store. As indicated in Fig. 1, if the resident storage medium is not core store, programs are paged into core prior to execution. The basis for establishing the resident storage medium for a particular program is determined by response time and frequency of execution requirements. For example, call-handling programs are core resident, diagnostic tests are disk resident, and system update programs are tape resident. Core-resident programs are backed up on both disk and tape: disk-resident programs are backed up on tape; and tape-resident programs have no on-line backup.

The No. 4 ESS software is packaged into approximately 800 individual PIDENTs, or loadable modules, ranging in size from several hundred to several thousand instructions. Each PIDENT generally implements one of many tasks required to perform an entire system function. An example of a PIDENT is the program that performs network connections, one task required as part of the call-handling function.

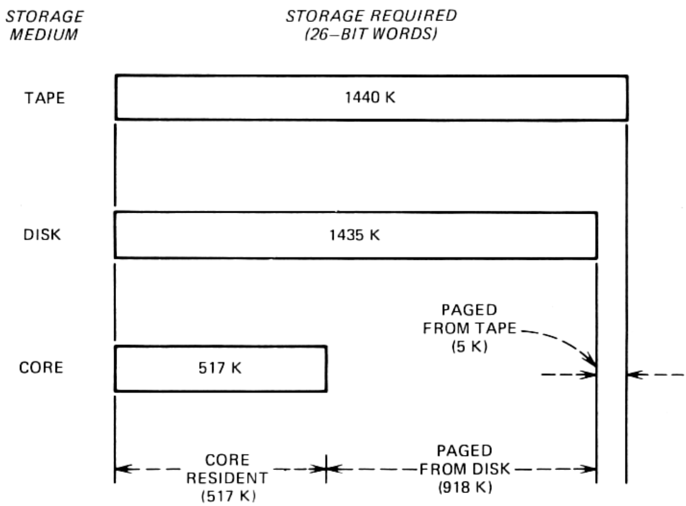

Fig. 1—Program storage media.

2.1 Operating system

The No. 4 ESS operating system can be conceptualized as a structure in three layers: executive control and audits, maintenance interrupts, and system integrity. Each layer is further subdivided into a number of activity levels that establish task priority within a layer. This structure is similar to that used in No. 1 ESS.[5] The principal difference is that timed interrupts are not used to perform high-priority operational functions. These functions are instead performed by using the interject approach. Figure 2 shows the relationship between the three layers of the operating system and, in decreasing order of priority, lists the 14 activity levels provided.

Executive control. The lowest layer of the operating system is the executive control, which performs job scheduling and sequencing in a normal, fault-free environment. Executive control schedules two levels of system activity: base level and interject level. As shown in Fig. 3, base level is a simple scheduling loop; the last task executed is succeeded in the scheduling order by the first. Base level tasks receive equal scheduling priority with each task served once per base. The time to complete one cycle of the base level schedule, called the "base-level cycle," varies with the instantaneous load on the system. It typically ranges from 11 milliseconds (ms) with no traffic load, to about 35 ms at a traffic load of 500,000 calls per hour. Examples of base-level tasks are the processing of new trunk service requests, administrative tasks (such as traffic measurements and network management), and frame diagnostic activity. Base-level programs are designed to divide their processing into a maximum of 3-ms segments, returning to executive control at the end

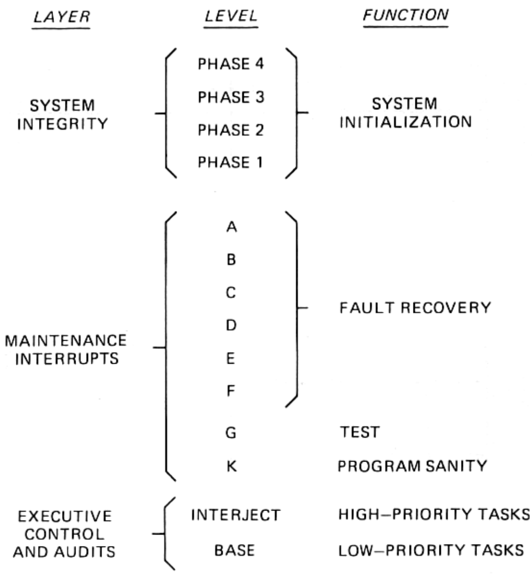

Fig. 2—Program activity levels.

of each segment. A task will be reentered on subsequent base-level cycles until completed. This allocation of 3 ms per task per base cycle is required to meet cross-office signaling delay requirements and prevents a single task from dominating system resources. This approach also helps limit the variance of the base-level cycle, which is important from the standpoint of designing stable real-time overload detection and control mechanisms.

Interject level is a scheduled programmed interruption of base-level activity, nominally every 10 ms, to do high-priority system tasks such as critical call-handling timing. As illustrated in Fig. 3, when each base-level task returns to executive control after completing 3 ms of processing, a check is made to determine whether 10 ms has elapsed since the last interject was serviced. If so, a transfer is made to the interject class of program tasks, which are linked together serially. At the conclusion of the final interject-level task, executive control returns to base-level processing. Since interject is a planned interruption of base activity, no extra overhead is required for saving and restoring processor registers. Also, the possibility of memory interwrite problems with other program activity levels is reduced when compared with a timed-interrupt approach.[5]

Maintenance interrupts. The second layer of the operating system provides detection and correction capabilities for hardware faults handled by maintenance interrupts. Maintenance interrupts (A through F) are triggered by the failure of some hardware fault detector (a parity

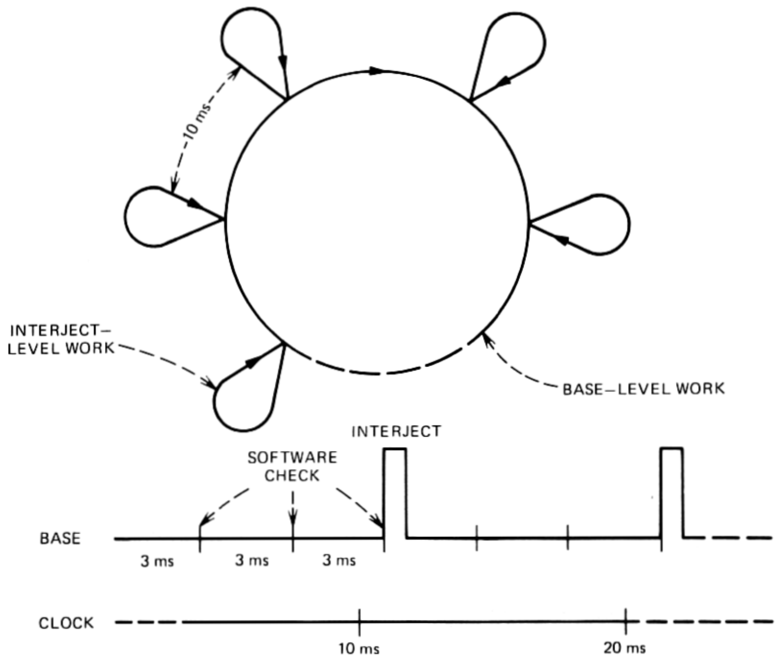

Fig. 3—Executive control.

check circuit, for example), and force the transfer or program control to the appropriate fault recovery program.[6] While the fault-recovery program is actively identifying and switching the faulty unit out of service, all other system operations are suspended. Fault-recovery programs are not segmented and will continue to completion unless a higher-level interrupt is triggered.

At the conclusion of the recovery action, normal system processing is restarted either at the point interrupted or at a "safe" starting point in the base-level cycle. Maintenance interrupts A through E are triggered in response to fault detectors resident in the 1A Processor. The F-level maintenance interrupt is reserved for units on the peripheral bus connected to the 1A Processor. These units include the network frames, signal processor, network clock, CCIS frame, and several 1A Processor units. G-level activity is restricted to such system troubleshooting tasks as obtaining data on complex field problems, and must be manually activated. K level provides an overall sanity check on the serving of 10-ms interject requests.

System integrity. The third and highest layer of the operating system, system integrity, provides a hierarchical structure of system initialization in four phases.[6] Although any phase may be manually selected and executed, automatic entry into the initialization sequence generally begins at phase 1 and progresses in sequence, if required, up to phase 3. Phase 4 can be activated only by manual request. Each phase incorporates additional initialization measures and, as a consequence, further impacts calls. Phase 1 has no effect on call processing. Phases 2 and 3 maintain connections on stable calls but tear down calls which are actively being processed. Phase 4 initializes all calls in the system to an idle state.

2.2 Operational features

The operational features of No. 4 ESS consist of call-handling capability for three signaling types: Multifrequency (MF), Dial Pulse (DP), and Common Channel Interoffice Signaling (CCIS); data administration software to modify the system translation data base; and surveillance software such as network management and traffic and plant measurements to monitor the efficiency of the No. 4 ESS system in switching calls.

Design of the call-handling software reflects heavy emphasis on efficiency, defensiveness, and ease of modification. The use of a single very large call register and a two-word trunk register, elimanition of peripheral order time buffering, per-call event consistency checks on data, and extensive employment of high-level macros are a few of the techniques used to meet the call-handling design objectives. Section III of this paper provides further detail on the call-handling design.

The primary requirement on the design of the data administration software[7] was to provide a simple, reliable man-machine interface for altering the system data base. The use of a set of system-generated forms, displayed on a CRT, provides a highly reliable and easy-to-use "fill-in-the-form" approach to data administration and permits the use of straightforward English text in the process.

The design of the surveillance software for No. 4 ESS is characterized by interpretative-table-driven program designs that permit the easy alteration of the output formatting and data processing algorithms of the network management and traffic and plant measurement functions.[8] This technique allows new surveillance features to be provided by a process similar to updating the system data base and is not tied to new program issues.

2.3 Maintenance features

The maintenance features of No. 4 ESS consist of trunk maintenance capabilities including manual and automatic trunk testing and frame maintenance provisions such as frame diagnostics and highly automatic trouble location procedures.

The trunk maintenance software system[7] controls the actions required for manual trunk tests via the 51A test position and automatic trunk tests via code line calls. In addition, trunk maintenance software automatically collects a wealth of data on ineffective machine attempts in the office and performs threshold analysis on this data to identify marginal trunks not uncovered during routine trunk testing.

The diagnostic program design for No. 4 ESS is characterized by a highly structured design approach which consists of a table of diagnostic tests and their expected results, specified by macros for each frame type, and a general interpretative control program.[6] The tests are disk resident and are paged into core for execution. Unlike previous ESS designs, several diagnostics can be active at a given time and the maintenance-crafts force receives the identity of the packs to be replaced at the conclusion of the diagnostic instead of receiving trouble number requiring manual lookup.

III. CALL HANDLING

3.1 Introduction

The call-handling programs control a call from origination to disconnection. These programs, either directly or indirectly, cause the performance of all the logical operations required to connect a call on an incoming trunk to the proper outgoing trunk, to supervise the connection, and to eventually disconnect the two trunks.

This section describes first the designs objectives for the call-handling programs and the capabilities of the No. 4 ESS. Next the hardware, data structures, and program structures that support call handling are described. The final subsection describes the operation of the call-handling program with a functional description.

3.1.1 Design objectives

The call-handling programs are designed to handle the three types of signaling provided by No. 4 ESS: MF, DP, and CCIS. Also, a means is provided for the collection of charging information on toll calls utilizing a Centralized Automatic Message Accounting (CAMA) feature. The call-handling design assumes that the No. 4 ESS will communicate only with other switching machines or operators but not directly with customer station equipment. Also, the initial No. 4 ESS call-handling design does not provide any special features required for gateway office operation or switching of international traffic.

In setting up the call-handling design for No. 4 ESS, a number of design objectives were established.

(*i*) A No. 4 ESS must meet specified traffic capacity objectives. The minimum required traffic capacities for No. 4 ESS are given in terms of switched attempts per hour, network CCS per hour, and trunk terminations. *Attempt processing:* A switched attempt comprises recognition of an incoming call, establishment of a connection through the network, and control of the associated signaling functions. At a load corresponding to peak capacity, other normal machine functions must also be performed (e.g., maintenance, traffic administration, network management). False attempts (i.e., incoming trunk seizures with no digits pulsed), within the limits shown, should not reduce the switched attempt capacity. Minimum requirements are as follows: engineered switched attempt capacity—550,000/hr (based on 10 high day busy hour average 500,000, 10 percent peak day increase 50,000); false attempts—66,000/hr. *Switching network load:* Network load capacity is a function of the number of connections through the network and their holding times, averaged over an hour. It is measured in CCS/hr (or in erlangs) and must be related to a probability of blocking of new calls to be meaningful. Two load levels, corresponding to the design objective for a maximum engineered load and for a peak load, are most useful. Minimum basic requirements are as follows: engineered load (0.5 percent first trial matching loss), 1,000,000 CCS/hr; peak load (10 percent first trial matching loss), 1,700,000 CCS/hr. *Network terminations:* Network terminations are the connecting points for the incoming, outgoing, and two-way trunk circuits. The two-way trunk circuits are considered to have a single network-connecting point. The network also terminates service circuits, tone and announcement circuits, and test circuits. Minimum basic requirement: 107,000 terminations (includes service circuits).

(*ii*) It is of only slightly lower priority that No. 4 ESS meet certain specified performance objectives at its rated capacity. Table I gives the mean time for several (but not all) of these objectives.

(*iii*) The system performance under overload must be reasonable. Absolute throughput should not be degraded even if the offered load exceeds the system capacity. Although performance requirements are relaxed somewhat in overload, it is anticipated that the system capacity will be limited by the tightness of system performance requirements rather than real-time exhaustion.

(*iv*) It is important that the system have a high degree of reliability. Both hardware and software faults can affect reliability. Of concern here are the techniques that keep the call-processing system operative in the face of such errors.

(*v*) An important characteristic of a large call-processing system is simplicity. Simplicity is a difficult concept to define, since it takes different forms in different instances. For example, the utilization of subroutines can lead to simplicity since a subroutine to perform a function can be written and debugged once and then used by other programs. On the other hand, straight-line coding can lead to simplicity because one

Table I—Performance objectives

Performance Measurement	Mean Time
Seizure time [seizure on incoming trunk (ICT) to transmittal of beginning of wink]	60 ms
Address time [receipt of last translatable digit to seizure of outgoing trunk (OGT)]	100 ms
Network path closure time (end of outpulsing to ICT-OGT connection)	100 ms
Cross-office answer time (receipt of answer on OGT to transmittal on ICT)	22 ms
Immediate-start time (seizure on ICT to ready to receive first digit)	50 ms
Response time (receipt of wink on OGT to transmittal of first address digit)	100 ms

does not have to keep track of transfers, data shifts, interfaces, etc. In fact, simplicity of a system may well come through the language in which it is written. One reason that simplicity is so important is that many other desirable properties follow from it: a simple system is easier to design, debug, make changes in, expand, etc.

(*vi*) Traditional objectives, such as efficient memory utilization and cost, have been considered, but the first concern has been given to the above considerations.

3.1.2 Signaling types

The address information (digits) that identifies the destination of a particular call and is transmitted between offices is classified according to the method of sending the information from one office to another. As mentioned previously, three types of signaling will be provided by the No. 4 ESS: Multifrequency (MF) pulsing, Dial Pulsing (DP), and Common Channel Interoffice Signaling (CCIS).

Multifrequency (MF). The No. 4 ESS is capable of sending MF digits to another office at either of two rates, one MF digit every 140 ms or one MF digit every 100 ms. These two pulsing rates are commonly called 7 and 10 pulses per second (pps), respectively. The No. 4 ESS also accepts MF pulses at either of the two rates. Also, No. 4 ESS can delay 0, 20, 80, or 220 ms between the receipt of the start-pulsing signal and the start of MF outpulsing.

Dial Pulse (DP). The No. 4 ESS accepts and transmits dial pulses at a nominal 10 pulses per second. No. 4 ESS will delay 280 ms between the seizure and start of DP outpulsing on DP immediate-start trunks. On non-immediate-start trunks, a 70-ms delay exists between the receipt of the start-pulsing signal and the start of DP outpulsing.

On DP immediate-start incoming trunks, the No. 4 ESS delays 90 ms between the seizure and start of digit collection to ensure that no digits have been missed. Any change of state during this 90-ms interval will be detected and cause the call to be handled as an ineffective attempt.

Table II — No. 4 ESS trunk types

Category	Trunk Types
Incoming intertoll, incoming toll connecting	DDD access, CAMA, TSPS, secondary intertoll, intertoll
Outgoing intertoll, outgoing toll connecting	Toll completing, intertoll, secondary intertoll, INWATS, rate and route operator, rate-quote operator
Two-way intertoll, two-way toll connecting	MF-MF intertoll, MF-MF toll connecting, DP-DP intertoll, MF-DP intertoll, CCIS-CCIS intertoll
Local tandem one-way	Tandem, tandem completing, intertandem completing, intertandem incoming, intertandem outgoing
Local tandem two-way	MF-MF intertandem, MF-MF tandem connecting

Common Channel Interoffice Signaling. No. 4 ESS, as part of its original design, provides CCIS.[9] With CCIS all address and supervisory signals are transmitted over a CCIS signaling network to which all toll switching machines in the Bell System will eventually be connected. No. 4 ESS provides the full complement of CCIS features available to the intertoll network.

3.1.3 Trunk types

The No. 4 ESS handles several types of incoming and outgoing intertoll, toll connecting, and local tandem trunks. These trunk types are listed in Table II. The No. 4 ESS also interfaces with 3CL-type switchboards and is capable of operation with the Traffic Service Position System (TSPS).

Service circuits. No. 4 ESS has a variety of service circuits necessary to complete call-processing functions. All service circuits are treated as one-way outgoing trunks and, therefore, cannot originate service requests to the system. Service circuits provided in No. 4 ESS include: tones (e.g., reorder), announcements (e.g., vacant code announcement), MF receivers, MF transmitters, and CCIS continuity check transceivers.

CAMA (Centralized Automatic Message Accounting). The No. 4 ESS can be equipped with up to 72 CAMA position trunks. These trunks are used to connect incoming calls to a CAMA operator when it is necessary for the operator to obtain the calling-party directory number directly from the customer. This procedure is identified as Operator Number Identification (ONI) service.

The No. 4 ESS is compatible with those switching systems that have Automatic Number Identification (ANI) equipment. In this arrangement, the ANI equipment identifies the calling party and forwards the calling party identity via MF pulsing to the No. 4 ESS office where the charging information is accumulated. In the event of ANI failure or inconsistent calling-party charge data, the No. 4 ESS uses ONI to identify the calling party.

3.1.4 Trunk characteristics

On-hook-when-idle operation. The No. 4 ESS software is designed to interface with on-hook-when-idle trunks. Any trunks requiring off-hook-when-idle operation for a connecting office will require a special trunk circuit at the No. 4 ESS office to convert off-hook-when-idle to on-hook-when-idle signaling.

Ring-forward signals. A ring forward is a timed on-hook followed by an off-hook signal on the ICT used by an operator to reestablish contact with another operator. Its duration must be greater than 60 ms but less than 200 ms to be a ring forward, with a nominal value of 100 ms, and it is allowed only on specified trunks.

There are two types of ring forward: an M-lead wink or a simplex ring forward, which is a 130-volt ac signal. Although No. 4 ESS can send either type, it can receive only the M-lead wink. Incoming trunks using the simplex version of the ring-forward signal have a special trunk circuit at the No. 4 ESS to convert it to an acceptable on-hook wink.

Wink start, delay dial-start dial, and immediate-start operation. No. 4 ESS will handle Wink-Start (WS) and Delay Dial-Start Dial (DDSD) signaling on MF incoming and outgoing trunks. Wink start, delay dial-start dial, and immediate-start signaling will be handled for DP incoming and outgoing trunks.

For both MF and DP incoming trunks with WS or DDSD signaling, No. 4 ESS will generate an off-hook wink with a 150-ms nominal value.

WS signals received by No. 4 ESS on MF or DP outgoing trunks must have a minimum off-hook interval of 100 ms and the on-hook must be received within 350 ms of the off-hook. In addition, the initial on- to off-hook transition must be received within 4, 5, or 10 seconds of the seizure, depending upon whether the trunk is intertoll, first-trial Toll Completing (TC), or second-trial TC.

MF DDSD signals received by No. 4 ESS must have a minimum off-hook interval of 100 ms, while DP DDSD signals must have a 60-ms minimum off-hook interval. For both cases, the delay dial signal must be received within 4 seconds of the seizure. The start dial signal must be received within 4, 5, or 10 seconds of the delay dial signal, depending upon whether the trunk is intertoll, first-trial TC, or second-trial TC.

STOP/GO operation. No. 4 ESS will handle STOP/GO signaling on DP outgoing trunks connected to a step-by-step office.

A STOP signal is an off-hook received on the outgoing trunk before the fourth DP digit has been outpulsed. This causes DP outpulsing to be suspended until the GO signal is received. The GO signal is an on-hook that must be received within 4 seconds of the STOP. DP outpulsing will resume 210 ms after receipt of the GO signal.

Echo suppressors. No. 4 ESS will work with trunks equipped with no echo suppressor, half echo suppressor, or full echo suppressor. Software control of echo suppressors is provided via distribution points.

Glare. Glare exists when both connecting offices simultaneously seize the same two-way trunk for use as the outgoing trunk in a call. On MF and DP WS and DDSD trunks, glare will be detected by the failure to receive the WS or DDSD signal in the specified time. Glare cannot be detected on DP immediate-start trunks. No. 4 ESS can detect and resolve glare conditions on nonimmediate-start and CCIS two-way trunks.

When glare is detected, two courses of action are available to the No. 4 ESS offices: it can (*i*) release the trunk to the incoming call or (*ii*) continue sending its connect signal until the other office recognizes glare and returns a start dial signal. When glare is detected on a two-way operator trunk, No. 4 ESS will always take action (*i*) above. For all other two-way trunks, one of the two connecting offices will be designated as the control office. For this case, the control office will take action (*ii*) above while the noncontrol office will take action (*i*).

CCIS continuity check. CCIS does not use the voice-transmission path to pass signaling information; therefore, continuity of the voice path must be checked separately. A CCIS continuity check transceiver is connected to the outgoing trunk and a 2010-Hz tone is sent over the trunk. The connecting office loops the tone back to the transceiver where it is subsequently detected. A continuity signal is then forwarded to acknowledge a successful CCIS continuity check.

3.2 Call-handling system organization

3.2.1 Hardware interfaces

The peripheral hardware (primarily the signal processor and time-division network) used in the No. 4 ESS had a significant effect on the call-processing design. In addition, since the periphery operates at a speed compatible with the processor, programs can communicate with the periphery in real time without the use of buffers. In previous ESS systems where the periphery was relatively slow compared with the processor, it was necessary for programs to time-buffer orders to the periphery. That is, a program would have to seize a buffer (possibly having to queue if all buffers were in use), load it with the necessary orders, activate the buffer, and then wait until all actions in the buffer were completed.

Signal Processor (SP). The SP is a No. 4 ESS peripheral unit that performs signaling and supervisory functions for up to 4096 E&M-type trunks. There can be up to 24 SPs in a No. 4 ESS office. The SP contains

6 bits of memory (called T bits) for each of its associated trunks. This memory is set by the call-handling programs and used by the SP to control its per-trunk functions. The SP scans each of its trunks once every 10 ms to determine if a change of state on its E lead has occurred. When a change of state does occur, the SP will make an entry into one of its output buffers if the trunk's T bits are set to a supervisory scanning state. The trunk's T bits can also be set to a state which will cause the SP to perform 30- to 40-ms hit-timing before reporting a state change.

The SP has four output buffers for reporting trunk status. The particular buffer used is a function of the type of report and the state of the trunk's T bits. One buffer is used for all high-priority reports (e.g., answers) and is interrogated by interject-level programs nominally every 10 ms. The other three buffers are interrogated at a slower, nonfixed rate by base-level programs. One buffer is used to hold only reports of new originations, another is used primarily for reports associated with digit functions, and the third is used primarily for reports associated with trunk abandons.

The SP also performs actions associated with digit reception and outpulsing for MF and DP trunks. Each SP can have up to 32 MF receivers connected to it. The SP periodically scans each of its MF receivers to determine if a digit has been received. The SP can accumulate from one to four digits before making an entry in its output buffer. As with MF receivers, each SP can have up to 32 MF transmitters connected to it. The SP can hold from one to four digits for each transmitter and can send the digits at either 7 or 10 pps.

DP digit receivers are not used in No. 4 ESS. Rather, DP digits are determined by detecting and counting changes of state on the trunk. The SP performs this function when such action is indicated by the trunk's T-bit state. The SP will make an output buffer entry after each DP digit is received. The SP will also detect and report abandons that occur during digit reception. The SP performs the reverse operations for DP digit outpulsing. The SP holds and transmits one digit at a time on the DP trunk. The SP also delays for the proper interdigital time before requesting another digit.

The SP can be used to report abandons on a trunk rather than simple changes of state from off-hook to on-hook. When indicated by a trunk's T-bit state, the SP will report an adandon when a change of state from off-hook to on-hook lasts more than 180 ms. In this state all other changes of state are ignored by the SP.

Time-division network. The No. 4 ESS switches all calls through a time-division switching network. This network switches data in Pulse Code Modulated (PCM) form with parity. Continuity of a path through the network is assured by the continued reception of good parity on the talking path. This parity bit also allows the detection of false crosses within the network because crossed paths will cause mutilation of the parity bit. Therefore no special continuity check of the cross-office talking path by program is required. In addition, since the switched portion of the network carries PCM code and not voltage levels directly, there is no need for a cut-through relay to protect the network. The design of the time-division network allows all network operation to at electronic speeds occur compatible with the 1A Processor.

CCIS terminal. The CCIS signaling system requires a fully synchronized bidirectional data link passing data at 2400 bits per second. The CCIS terminal is a programmed controller which administers the routine tasks necessary to maintain synchronization of the data stream with the other office. The terminal recognizes incoming signaling messages, separates them by priority, and buffers them for the 1A Processor. Messages containing errors are intercepted by the terminal and a request for retransmission from the other office is made. The terminal also accepts messages from the 1A Processor and queues them for transmission on the data link by priority. It also maintains a history of transmitted messages so they may be retransmitted upon request from the other office. Thus the 1A Processor has a very simple software interface to the data link which takes very little real time to administer.

3.2.2 Data structures

There are various data structures used by call-handling programs. These data structures include memory facilities, the call event reporting structure, link lists, queues, and the software timing structure. The two major memory facilities are the call register and trunk register.

Call register. A Call Register (CR) is a 64-word block of call store memory that is used for temporary storage of information during call setup. The CR is large enough to contain all information needed for processing a call. Thus, additional blocks of information need not be linked to the main register to obtain more storage space. Furthermore, any information that is derived and may be required later will be stored in the CR.

CRs are not dedicated on a per-trunk basis. Instead, there is an engineered number of CRs per office. Idle CRs are link-listed to minimize the time required both to find an idle CR and to restore a CR to idle.

Trunk register. Trunk Registers (TRs) are two-word blocks of call-store memory assigned on a per-trunk basis. TRs contain dynamic information about the current state of the trunk. All TRs are in a continuous block of call-store memory.

Reporting structures. When an internally or externally generated call stimulus (report) is detected, control is transferred to a processing routine that is identified by the type of report and the state of the call.

When a CR is not associated with a call, the state of the call is specified by the TR state code. In this case, the particular processing routine is normally determined by using the TR state code to index a table associated with a particular type of report.

When a CR is associated with a call (indicated by the TR state code), the state vector (SV) in the CR is used to specify the state of the call. The SV identifies the state of the call and is made up of eight components. Each of the eight components represents what is, in general, an independent function or part of the call. Hence, each component part of the total state vector can change independently without affecting the other state vector components. This structural arrangement simplifies the program design because it reduces the interaction of the various parts of the program and saves real time during the processing of a call.

When a report associated with one of the SV components occurs, the task program will combine the existing SV component value and the type of report and use the resultant value to determine the address in the program for processing the report. Normally, all reports can be processed independently and without checking the values of other SV component values.

Link and engineered lists. Two basic types of lists of software structures are provided in No. 4 ESS, two-way link lists and engineered lists. These lists are used to link together related software facilities, to provide timing functions, etc. For example, all idle call registers are placed on an idle-link list.

An entry on a two-way list has a linkage to the previous entry and another linkage to the succeeding entry on the list. Two-way lists are used in No. 4 ESS to minimize the overhead associated with entering and removing entries from the list. They can be of any length and are also easily searched for consistency, for example, by an audit routine.

Engineered lists, on the other hand, are used to provide special functions, for example, timing. These lists have a specific (engineered) number of entries.

Queues. A queue is a group of facilities of a given type which are all waiting for a second type of facility to become available for use. In the No. 4 ESS call-handling program, the only structures which can be on a queue for a facility are CRs and TRs. All queues consist of link-listed CRs and TRs. This provides a simple method of giving first-in, first-out service to the entries on a queue. CR and TR queues are two-way link lists.

The task program administering each queue is entered periodically to determine if there is an entry on the queue and, if so, if there is an idle facility available. If both conditions are met, then the first entry is removed from the queue and a transfer is made to the program needing the facility.

Software timing. Various types of timing are required during a call. For example, permanent signal–partial dial (PSPD) timing is required on all MF incoming trunks and disconnect timing is required when an on-hook report is received on the incoming trunk for a call in the talking state. In the first case, a CR is associated with the call while in the second case, only TRs are associated with the call. The No. 4 ESS call-handling programs use four types of timing: CR timing, TR timing, dedicated word timing and engineered list timing. CR and TR timing are performed by linking the particular structure (CR or TR) to a timing list. Both interject-level (accurate to within 10 ms) and base-level (accurate to within 1 base-level cycle length) timing can be performed with the CR and TR.

Dedicated word timing is another method of performing timing using the CR associated with a call. One word in the CR is reserved for timing and contains an active flag, index, and time-out time. The active flag, when set, indicates that timing is active in this CR. The index specifies which type of timing is being done, and the time-out time specifies when a time-out will occur.

When the dedicated word-timing task is scheduled, each CR in the office is interrogated for time-out if the active flag is set. This method of timing is used only where relatively long, inaccurate, and universal timing is required. An example of this is PSPD timing.

The fourth method of timing is performed with the use of an engineered list. This method of timing is used when a CR is not associated with a call and when the TR cannot be added to a timing link list. Two engineered lists are provided, one for interject- and one for base-level timing.

3.2.3 Call-handling program structure

The call-handling programs are structured in a three-level hierarchy as shown in Fig. 4. The *task dispensers,* which are entered directly from executive control, interface with the signaling hardware (signal processors and CCIS terminals). The task dispensers pass call-handling stimuli to *task programs.* While performing call actions, the task programs may use one or more *call-handling subroutines* to execute repetitive or highly specialized actions. The task programs also interface with other operational programs (e.g., translations, trunk maintenance) as described in Section 3.2.4.

Task dispensers. The call-handling task dispensers are responsible for distributing call-related stimuli to the call-handling task programs. The stimuli can be either external (from the signaling hardware) or internal (timing or queuing reports). The task dispensers operate on both base and interject level, dispensing high- and low-priority reports re-

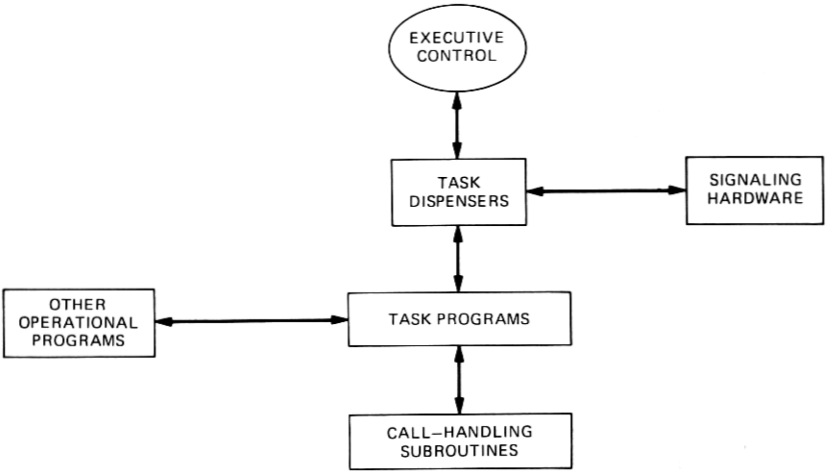

Fig. 4—Call-handling structure.

spectively. There are two basic task dispenser programs, the MF/DP task dispenser and the CCIS task dispenser.

The MF/DP task dispenser program interfaces directly with the signal processor. On each entry from executive control (interject or base), the signal processor buffers are examined for call-relevant reports (e.g., on-hooks, off-hooks, digits, stop dial). These are dispensed sequentially to the task programs for processing. The MF/DP task dispenser also dispenses internally generated time-out conditions to the task programs. Again, during each entry from executive control, the timing and queuing lists are examined for time-out conditions. If a time-out condition exists, the appropriate task program is entered to process the particular stimuli. The task dispenser remains in control until all relevant internal and external stimuli are processed or until an overload threshold is reached. The overload threshold provides a control on the amount of activity processed by the system during any base cycle.

The CCIS task dispenser interfaces with the CCIS terminals and, like the MF/DP dispenser, polls the terminal buffers for CCIS messages. If messages are present they are dispensed sequentially to the appropriate CCIS task program.

Task programs. Call-handling task programs, in general, are used to perform the specific actions that switch calls. In No. 4 ESS there are two types of task programs, MF/DP task programs, which handle MF and DP calls, and CCIS task programs, which handle CCIS calls. The task programs are entered from the task dispensers in response to a particular call stimulus. The task program will investigate the present state of the call by examining the TR and CR state codes. Depending on the present state of a call and the new call stimuli, the appropriate actions to advance the call are executed. For example, if a call is in the "waiting-for answer" state and an off-hook stimulus is received by the task program on the outgoing trunk, the task program will verify the validity of the report, transmit an answer condition on the incoming trunk (e.g., off-hook), and advance the call state to "talking."

Call-handling subroutines. Certain repetitive or specialized call-handling functions in No. 4 ESS are designed as subroutines where they may be accessed by several task programs. Examples of subroutine actions are: the seizing and initializing of a call register, the connection of incoming trunk to outgoing trunk, the hunting of a service circuit and the pegging of a traffic counter.

3.2.4 Call-handling software interfaces

As illustrated in Fig. 4, the call-handling programs interface with other operational programs during the processing of a call. These interfaces were established to allow independent software development of major operational functions such as audits, translations, and network management. Where these functions overlap during the processing of a call, clearly defined interfaces were established.

Audits. The call-handling programs make defensive checks to help ensure that data associated with a call has not been mutilated. When an error is found, an audit program is called. In general, when an audit program of this type is entered, the affected data structures (e.g., CR, TR) is isolated, information for a teletypewriter printout is generated, flags are set to demand additional audits to restore the structures on a deferred basis, and the affected call is terminated.

Translations. During the processing of a call, translation data is needed to complete the call. This data includes the trunk identification, trunk signaling characteristics, routing, and digit prefix and delete information. The call-handling programs call the appropriate translation retrieval routines to obtain the needed data.

Trunk maintenance. Trunk maintenance programs are called by the call-handling programs whenever a possible trunk-related hardware problem is encountered. For example, if not enough MF digits are received in the allotted time, the MF receiver used on the call is passed to trunk maintenance for testing on a deferred basis, since it may not be able to recognize digits. The incoming trunk is also passed to trunk maintenance for possible testing. The trunk maintenance programs are also entered to handle various test calls.

Network management. The call-handling programs interface with network management programs to modify the routing of calls based on network management actions. Network management routines also maintain a data base of call-completion statistics for later analysis.

Overload Control. The overload control program interfaces with the

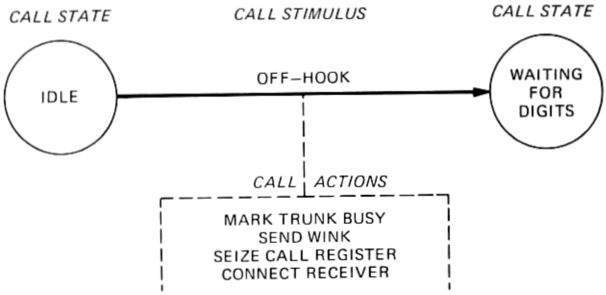

Fig. 5—Call-handling state control.

call-handling programs to control the number of originations processed in any one base-level cycle. In this way, strict control is maintained over the length of the base-level cycle.

3.3 Call-handling functional description

The handling of calls in No. 4 ESS follows a stimulus-response arrangement as illustrated in Fig. 5. Each call is defined in terms of a *call state* and each call stimulus (e.g., off-hook) advances the call state by performing specified call actions. In Fig. 5 the call state is advanced from idle to waiting-for-digits by the call actions of responding to the seizure, seizing a CR, and connecting the ICT to a receiver. Each stimulus is independently brought into the system by the task dispenser and passed to a task program for action. The call state is maintained in the trunk register and call register if attached.

The remainder of this section describes the No. 4 ESS call-handling programs in terms of their signaling capabilities and major functional modules. The handling of MF, DP, and CAMA calls is described in detail while the handling of CCIS is described in abbreviated form. The major call-handling subroutines of digit reception, final handling, and network actions are also described.

3.3.1 MF and DP Signaling

No. 4 ESS recognizes originations on MF DDSD (Delay Dial—Start Dial), MF WS (Wink-Start), DP DDSD, and DP immediate-start incoming trunks via the signal processor. When the origination is on an MF trunk, a call register (CR) is seized and linked to the trunk register (TR) associated with the trunk and an MF receiver is connected to the incoming trunk (ICT). An MF origination will queue if there is not a CR and a receiver available. A DP origination does not have a CR associated with it until after some of the digits have been received in order to reduce the CR holding time. Overload can limit the number of MF originations served per base level and the number of DP originations served per interject. An origination is put on queue if an overload threshold has been exceeded.

When an origination has been accepted, a wink must be sent to the preceding office if the trunk is either DDSD or WS. On MF DDSD trunks, the trunk circuit at the other end must receive the off-hook of the wink within a specified time or the call will be aborted. Therefore, if an origination on an MF DDSD trunk is queued for any reason, the off-hook is sent when the queuing starts.

When the "handshaking" has been completed, the digits of the called number are collected. When the ICT is DP DDSD, No. 4 ESS blinds itself for 40 ms after the start dial has been sent so that it does not interpret reflections from the start dial as dial pulses. When the ICT is DP immediate-start, the No. 4 ESS does early digit timing to ensure that the first part of a digit has not been missed because there is no delay before the digits are outpulsed to No. 4 ESS. If there is an SF (Signal Frequency) set involved, the timing is for 90 ms; if there is no SF set, only 60-ms timing (+30 ms in the signal processor) is performed. Any change of state during this timing is interpreted as an indication that part of a digit may have been lost.

For those trunks with an SF set, the SF set provides the hit-timing. When there is no SF set, the SP (Signal Processor) provides the hit timing.

The digits of the called number are used to obtain the routing information, which resides in the memory of the No. 4 ESS (see Section 3.3.4 for details of the digit reception function). An outgoing trunk is selected from the list of trunks supplied by the routing data and a path between the incoming and outgoing trunks is selected and reserved. If the outgoing trunk (OGT) is MF, the availability of a transmitter is checked, although it will not be connected until later.

A preliminary glare check on an OGT consists of performing a directed scan. If the trunk is off-hook, it is assumed to be in use by the other office and the No. 4 ESS backs off. On a one-way OGT, no glare is possible, so a glare detection is interpreted as a bad trunk. If glare is detected on a two-way operator trunk (which is really a one-way OGT), the No. 4 ESS backs off. The only true glare case can be encountered on the two-way OGT.

The appropriate handshake with the OGT is initiated. The OGTs handled by No. 4 ESS are MF DDSD, MF WS, DP DDSD, DP WS, DP immediate start, no-outpulsing, integrity check no-outpulsing [this is a WS no-outpulsing trunk for the five ACD (Automatic Call Distributor) feature], and two-way operator trunks. This last type requires that a miscellaneous point be operated to get access to the trunk.

When the handshake has been completed, the digits of the called number, after any necessary prefixing or deleting, are outpulsed. However, if the OGT is MF, the call first queues for an MF transmitter. When one has been obtained, it is connected to the OGT and outpulsing starts. If the connection fails, a new transmitter is obtained. A second failure aborts the call.

When outpulsing on a DP trunk, No. 4 ESS scans for a STOP (off-hook on the OGT) between digits. If a STOP is detected, outpulsing ceases and is not resumed until a GO signal is received.

When outpulses are sent on an MF OGT and the ICT is CCIS, the last digit to be outpulsed is held back until the COT (continuity) signals has been received. This means that there can be a delay between the tens and units digits.

Failures during MF outpulsing can be a result of a transmitter error, an off-hook on the OGT, or an on-hook on the ICT. This last case sends the call directly to final handling, and the OGT is idled as soon as the outpulsing of the current digit group has been completed. In the first two instances, the No. 4 ESS awaits the completion of the outpulsing of the current digit group before hunting a new trunk and trying again. In addition, if an off-hook occurs on a toll-completing trunk during the outpulsing, it will be considered an early answer; if this occurs on an intertoll trunk, however, it is considered to be an unexpected STOP and the call is aborted. A second transmitter error or a second unexpected STOP will abort the call.

When outpulsing has been completed, any transmitter is disconnected, the incoming trunk is connected to the outgoing trunk, and the CR is released.

If the ICT disconnects before an answer on the OGT is detected, the call is terminated and the connections are abandoned. If an answer is received, however, the call enters the talk state, in which state ring forward, clear back, reanswer, or disconnect can be received. Since No. 4 ESS employs calling-party hold, an on-hook on the OGT will be considered as a clear back. It will be passed on but will not be considered as a disconnect; only the ICT can disconnect. Once in the clear back state, an off-hook on the OGT is considered as a reanswer and is passed on.

Guard timing is the time given to the terminating office to idle a trunk. It prevents an attempt to use the trunk again before the terminating office equipment has had time to complete the release. It is always employed on an OGT and will also be used on an ICT if the M relay has been operated but digits have not yet been received. Guard timing is 1050 ms.

3.3.2 Common Channel Interoffice Signaling

There are two functional components of the operational software required for CCIS. The first component, CCIS link security, maintains the integrity of CCIS data links to other offices. The second component, CCIS call handling, handles all related signals associated with message circuits.

CCIS link security. The link security system administers the CCIS data links. These programs ensure that an optimum configuration is maintained for signaling facilities, that the error rate on the facilities is acceptable, and that alternate facilities are selected if a data link fails. Link security monitors and performs the synchronization protocol on the data link and provides a craft interface for facilitating repairs of the signaling link components.

CCIS call handling. The CCIS call-handling programs interface with signals received over the CCIS data links to switch calls between CCIS trunks and all other types of trunks. Initial CCIS incoming trunk actions handled by these programs include the analyzing of CCIS address messages, the connecting of a zero-loss loop via the time-division network to facilitate the interoffice continuity check performed on the CCIS incoming trunk, and the initiation of routing and outgoing trunk selection. Outgoing CCIS trunk actions include the control of the interoffice continuity check with a continuity check transceiver and the administration of backward failure signals. Once a call has reached the waiting-for-answer state, the CCIS programs will administer answer, ring forward, and disconnect signals.

An important factor in the CCIS program design results from the error-correcting characteristic of the signaling link. Since signals transmitted on the data link may contain errors and be retransmitted, it is possible to receive CCIS call signals out of sequence. Although this happens infrequently, the CCIS call-handling programs must process these out-of-sequence signals to successfully handle affected calls.

3.3.3 CAMA

The No. 4 ESS can provide CAMA (Centralized Automatic Message Accounting) service for all customers in class 5 offices which home on it and which do not have LAMA (Local Automatic Message Accounting), for multiparty line customers served by a LAMA office (since LAMA service is limited to one- and two-party lines).

All the data for a particular CAMA call are buffered in a dedicated Accounting Block (AB) and stored in a single entry on a nine-track 800-bpi AMA tape. The data consist of the calling number, the called number, the answer and disconnect times accurate to 0.1 second, plus any other information needed to correctly bill a call. The calling number may be reported to the No. 4 ESS via either ANI or ONI. In the ANI case, equipment in the local office where the call originates outpulses the calling number to the CAMA office. If the call is ONI or if there has been an ANI failure, the No. 4 ESS attaches an operator to the call to query the customer for the calling number. The operator then keys the calling number into the No. 4 ESS.

In addition to collecting and recording toll billing data, a No. 4 ESS office can be equipped to collect and record data on calls for which CAMA records are not usually made but which do come over CAMA trunks.

A No. 4 ESS CAMA office can handle a maximum of 8160 CAMA incoming trunks. No. 4 ESS can handle a maximum of eight NPAs (Numbering Plan Areas). It can interface with a maximum of 72 CAMA operator positions consisting of a keying trunk and a talking trunk. With these two trunks, the operator is able to key in the calling number over the keying trunk while receiving it from the calling subscriber over the talking trunk.

Signaling between No. 4 ESS and the CAMA positions can be on either a loop or an E&M signaling basis. If E&M signaling is used, a special trunk circuit allows the CAMA position to be a TSPS operator position. The CAMA operator may be at a regular CAMA cordless position, a cord switchboard modified for CAMA operation, or a TSPS No. 1 100B position. These positions can be in the same building with the No. 4 ESS, or they can be located remotely. To the toll office, however, they always appear to be at a remote location.

Incoming CAMA calls are processed by No. 4 ESS on a first-come, first-served basis. When a CAMA call requires ONI, the most idle CAMA position trunk is chosen to permit an even distribution of calls among CAMA position trunks. If no CAMA position trunk is available for a CAMA ONI call, the call is queued and the customer receives audible ring until a position becomes available.

A hardware system clock is used to accurately maintain the software time-of-day clock.

3.3.4 Digit reception

The digit reception module is responsible for performing all actions unique to the collection and analysis of digits received on an incoming call, including the calling number for a CAMA call. This module centralizes all call-handling actions relative to the accessing of translation and routing data, including interfacing with code restricting network management routines. The actions performed by this module are described below in terms of the digit and routing capability of No. 4 ESS.

No. 4 ESS will accept 3 through 11 digits (excluding the KP and ST digits) domestically via MF, DP, or CCIS signaling. It can translate on 3 through 9 digits to obtain the routing data. When outpulsing, it is possible to delete any number of digits and/or prefix up to 6 digits. The digit translation tables and the routing data blocks can be in core and/or disk storage.

When the incoming signaling is via MF, all digits are collected before a digit translation is requested, since the ST digit acts as a positive indication of end of dialing. When digits arrive via CCIS, they are all contained in one message, so when the digit translation is requested, all the digits have been received. When the incoming signaling is DP, however, there is no positive end-of-dialing indication. Five-second critical interdigital timing is performed after any digit which may have been the last, and 15-second noncritical interdigital timing is performed after any digit which is not expected to be the last. To avoid delaying each DP call by at least 5 seconds in determining end of dialing, the digit translation is overlapped with the DP digit collection. Normally, the routing information will be obtained before all digits have been received. Part of this routing information specifies the expected number of digits and, at least in the case of receiving the maximum number of expected digits, this information can be used to avoid the necessity of timing after the last digit.

Normal calls, maintenance calls, and internally generated test calls all use the same digit reception software to collect, analyze, and translate digits, although some calls, such as test calls, do not require the translation.

Any failure encountered in the digit reception process, such as an illegal digit (i.e., two consecutive KP digits) or a digit error (an MF code which is not 2-out-of-6), result in the call being sent to final handling for termination.

3.3.5 Final handling

The final handling module is called to idle facilities and update counters associated with calls that are not completed by the No. 4 ESS; i.e., Ineffective Attempts (IAs). An IA is any attempt recognized by the switching machine as a bid for service but which does not subsequently result in the call being completed in the desired manner. This category consists of the calls which do not reach the waiting-for-answer state for any reason. Once a call reaches the waiting-for-answer state it will be idled by normal call-handling actions.

When an IA is detected, there are facilities associated with that call which must be restored to an idle state so that other calls can use them. A "facility" is defined as a hardware facility, or a software structure: e.g., service circuit, trunk, call register, or timing-list entry. Final handling also connects the incoming trunk to a recorded announcement or tone as appropriate. It will time the announcement connection and will disconnect it if the incoming trunk does not disconnect within a reasonable amount of time. No. 4 ESS provides all necessary tones and announcements for toll service.

3.3.6 Network actions

The network-actions subroutine is the interface between the call-handling programs and the No. 4 ESS time-division network hardware, specifically the time-slot interchange (TSI) and the time-multiplexed switch (TMS). To provide this interface function, the network-action routines must hunt network paths, perform the path setup, maintain the residual path information necessary for releasing the path on demand, and must also administer the use of the network links to prevent crosses from occurring. Since the network-action routines are the only ones designed to access the time-division network, these routines must set up all the types of paths required by the call-handling programs. The types of routines provided include a normal 2-way path connect, a path reservation (used while one or both trunks are involved in another connection), monitor connections and a broadcast feature.

Path hunt strategy. The path hunt strategy chosen to select idle network paths must satisfy two criteria:

(*i*) It must not degrade the inherent blocking characteristics of the time-division network.

(*ii*) It must be real-time efficient; i.e., the run time of the hunt must be kept to a minimum.

In order to satisfy the first criterion, extensive simulations were made of the final strategy, proving that the network met its blocking objective with a substantial margin. To satisfy the second requirement, great care was exercised in the design of the data structures and the resultant program necessary to execute the network path-hunt algorithms. Once a path is selected between two endpoints, the network program causes that path to be set up in the network. In order for this function to be performed efficiently by the program, the network hardware was designed to accept data to set up paths in a form compatible with the internal data formats used by the network programs.

Network link administration. A map of the busy/idle status of each link on each time slot is maintained in the software. The purpose of this map is to ensure that only idle facilities are used in setting up new network paths. The map is composed of three parts. The first part, the time-slot map, keeps track of each of the 128 time slots to which a trunk has access. The second part, the A-link map, keeps status on the links joining the TSI to the TMS. Finally, the B-link map has the status of the links joining the two stages of switching within the TMS.

The two criteria which governed the design of these status maps were speed of access for the path hunt program and modularity to allow graceful growth of the map as the office grew.

The network map structure enables the network programs to keep track of link usage. It is also necessary to maintain information concerning active paths in the network, so they may be disconnected properly. Since this is trunk-related information, it is stored in the trunk register of one of the connected trunks. The trunk register also has a pointer word which identifies the other trunk which is connected via the network path. The two trunk identities and the path information contain enough information to identify all links used by the network path being held.

Monitor connections. The ability to monitor transmission for any trunk within No. 4 ESS is a requirement for trunk maintenance. It is possible to use the one-way transmission capability of the time-division network to connect a trunk circuit, which is already busy, to a special trunk maintenance bridging circuit. This allows measurement of transmission factors while the trunk is in use.

Broadcast of announcements. Connections to customers needing announcement or tone treatment are made by a special connect routine which allows the same announcement to be broadcast to many customers simultaneously. This is provided by dedicating a special TSI to this function, making a semipermanent network connection from the announcement source to the special TSI, and using the unique properties of the time-division network to make the announcement or tone available to as many as 1024 network ports.

IV. SUMMARY

This paper has described the operational software structure of No. 4 ESS with particular emphasis on the design of the call-handling software. The development of this software package was a large undertaking utilizing the skills of over 100 software designers. While some of the software philosophy used in No. 4 ESS was built on previous ESS systems, much of the design was totally new, particularly in the areas of handling and switching toll traffic through a digital switch and the accompanying administrative features. We acknowledge the effort of those designers, too numerous to mention, who contributed to the successful No. 4 ESS development.

REFERENCES

1. J. M. Scanlon and L. S. Tuomenoksa, "No. 4 ESS Software," Proc. Intl. Switching Symp., October 1976.
2. M. J. McPheters, S. Bloom, S. H. Tsiang, "Software Quality Control," Proc. 1973 Software Reliability Symp., April 1973.
3. J. S. Nowak, "No. 1A ESS—A New High Capacity Switching System," Intl. Switching Symp., October, 1976.
4. B.S.T.J. special issue on 1A Processor, *56,* No. 2 (February 1977).
5. J. A. Harr, E. S. Hoover, and R. B. Smith, "Organization of the No. 1 ESS Stored Program," B.S.T.J., *43,* No. 5 (September 1964), pp. 1923–1959.
6. M. N. Meyers, W. A. Routt, and K. W. Yoder, "Maintenance Software," B.S.T.J., this issue, pp. 1139–1168.
7. J. A. Giunta, S. F. Heath, J. T. Raleigh, and M. T. Smith, "Data and Trunk Administration and Maintenance," B.S.T.J., this issue, pp. 1203–1238.
8. T. V. Greene, D. G. Haenschke, and B. H. Hornbach, "Network Management and Traffic Administration," B.S.T.J., this issue, pp. 1169–1202.
9. B.S.T.J., special issue on CCIS (to be published).