

Common Channel Interoffice Signaling:

Signaling Network

By P. R. MILLER and R. E. WALLACE

(Manuscript received May 7, 1977)

The CCIS signaling network provides an efficient, reliable signaling capability for calls between CCIS-equipped switching offices. This article contains a description of the signaling network and the Signal Transfer Points. The discussion highlights the normal signaling message routing as well as the reaction to system failures.

I. NETWORK DESCRIPTION

1.1 General

The CCIS signaling network provides an efficient, reliable common channel signaling capability for all intertoll message circuits between CCIS-equipped switching offices. In order to provide cost-effective signaling for small, as well as large, trunk groups, the nonassociated mode of signaling¹ is used as the primary signaling mode. The system also has the capability for associated signaling in situations when trunk groups of sufficient size exist. The reliability of the signaling network is provided by redundancy in both signaling links and Signal Transfer Points (STPs). This redundancy is utilized by automatic routines at the network nodes which alter the routing of signaling information in order to bypass network failures.

1.2 Network configuration

The signaling network is composed of 20 STPs, two in each of the ten DDD regions, and signaling links interconnecting STPs and switching offices. The STPs in each region are redundant, in that each has the capacity to serve the entire region in the event of the mate STP failure. In the normal situation, signaling traffic is routed in such a manner as to balance the load carried by each STP in a region. Figure 1 shows two regions and the signaling links between various nodes. Signaling links




Fig. 1—CCIS signaling network.

which connect switching offices to the STP pair in the same region are called A-links. A-links are always provided in pairs, one link to each STP. The STPs in a given region are fully interconnected to the STPs in all other regions using signaling links designated as B-links. These links are always provided in groups of four, commonly referred to as a quad, between any two regions. Groups of signaling links, such as A-link pairs and quads which provide redundant signaling paths, are referred to as link complements. Under normal conditions the traffic load is evenly distributed over each link of a complement. Signaling links between the two STPs in a region are referred to as C-links. C-links do not carry signaling traffic under normal conditions but provide a signaling path during failure situations. Hence, the number of C-links provided is proportional to the total number of A-links. All of the C-links between two mate STPs are members of one link complement. In contrast to the nonassociated network, the signaling path between switching office C and office D in Fig. 1, referred to as an F-link, is an example of a link used in the associated signaling mode.

All signaling links are composed of a terminal and modem at each end connected by a Voice Frequency Link (VFL) as described in Ref. 1. A-links have the additional ability to operate with redundant VFLs. Figure 2 shows an A-link with duplicated VFLs, as well as the VFL test access features located at the STP.

1.3 Normal telephone message routing

Telephone signaling messages originate at a switching office and are routed through the signaling network, terminating at a distant switching office. The originating office usually begins this process with the identity




Fig. 2—STP VFL arrangement and dual VFL configuration.

of a specific trunk. This trunk identity is translated to a signaling link number and a label, the combination of which is uniquely assigned to the trunk. Messages pertaining to that trunk, composed of the label, message type, and possibly data, are transmitted to the STP over the designated link complement. The STP, however, does not use the entire trunk identity, since messages for trunks in the same trunk group have similar message routing. Instead the STP examines only a portion of the label, along with the incoming signaling link number, to determine the destination of the message. This portion of the label is called a band and consists of 16 consecutive labels. Trunk groups which consist of more than 16 trunks have additional bands assigned.

The STP, upon receipt of a telephone message, converts the incoming link number and incoming band to an outgoing link complement number and outgoing band by means of routing data stored at the STP. The message is then transmitted with the outgoing band number replacing the original band number on the outgoing link complement. The outgoing link may terminate at another STP (interregional trunk group) or at the terminating switching office (intraregional trunk group). In either case, when the message arrives at the terminating switching office, having traversed, under normal circumstances, no more than two STPs,

the incoming signaling link complement and label are translated to the specific trunk. At each node in the network the translation data are constructed to balance the traffic load between all links of a link complement and hence between each STP in a region.

Nontelephone messages do not contain labels and are processed at the receiving node. These messages are used for control, administration, and maintenance of the signaling network.

1.4 Network reaction to failures

The signaling links have two major failure symptoms. First, normal component failures occur, with mean times between occurrences measured in years and requiring manual repair activities. The repair interval is normally of a few hours or less in duration. Second, the VFL exhibits periods of high error rates caused by hits, fades, or other random phenomenon on the transmission facility. High error rates for short periods are processed by the error control mechanisms in the terminal as described in Ref. 1. If the duration of these periods is sufficiently long, (approximately $\frac{1}{3}$ of a second or greater), the signaling link is considered as possibly faulty and the signaling traffic temporarily rerouted. If the excessive error rate period persists for 3 minutes or greater a component failure is assumed and the network is reconfigured as explained below. Typically the time between VFL short interrupts is measured in fractions of a day and the duration of the interrupts is measured in seconds. Hence the network failure characteristics are summarized as short interrupts requiring no manual repair which occur orders of magnitude more frequently than component failures.

As an example, assume link A1 in Fig. 1 fails because of an excessive error rate. The failure would be recognized at both switching office A and STP-1. Office A would immediately send all traffic normally routed over link A1 to link A2. STP-1 would send all incoming signaling traffic destined for office A to STP-2 over the C-link complement. When telephone traffic is routed over a C link, an additional signal unit, called a header, is prefixed to the message. STP-2 would then route this traffic, using the information transmitted in the header, to office A over link A2. The terminals at both ends of link A1 would attempt to restore normal operation. If the link is restored within 3 minutes, the normal routing pattern is reestablished automatically and hence no alarms or other immediate notification are provided to the office personnel. Statistics on all failures are kept and reported periodically.

If the failure condition persists for longer than 3 minutes a component failure requiring repair is assumed to have occurred. When 3 minutes of outage has been recorded the terminals at both ends of the failed link are automatically removed from service, accompanied by the appropriate audible and visual alarms, and examined for faulty components. At this

time the network configuration is changed to provide optimal routing for the particular situation. In the example above, where link A1 is failed, STP-1 would notify the STPs in the other regions and all switching offices in the same region which route traffic to office A. This notification is in the form of special signal units, called Transfer Restrict (TFR) and Transfer Prohibit (TFP) messages, which are sent for each affected band. When notified, the STPs and switching offices send all traffic, which would normally traverse link A1, directly to STP-2, which then sends the traffic to office A on link A2. Thus, by configuring the traffic pattern for the affected bands to avoid C-links, signal unit delay is decreased and C-link capacity is reserved for short link interruptions. All other traffic to STP-1 would not be affected. This configuration would remain until link A1 is repaired and restored to service at which time the normal traffic pattern would be automatically restored. The algorithms at each node in the network are designed to handle all combinations of single and multiple link failures in order to route traffic as long as a signaling path exists. In the rare event that no signaling paths exist, the affected message trunks are automatically prevented from accepting new calls. These trunks would be immediately restored as soon as a signaling path is reestablished.

The network is also protected against an STP failure. In this event, each terminal at the failed STP autonomously notifies the distant end of the signal link. When notified, the STP in the distant regions, as well as all switching offices in the affected region, will immediately route all signaling traffic to the remaining STP.

Similar network reconfigurations take place whenever a signaling link is removed from service manually for facility rearrangement or other reasons. Manual removal always requires permission from the processors at both ends of the link. Permission granted implies that the network has been reconfigured to avoid this signaling link with no loss of signaling capability.

II. STP DESCRIPTION

The STP basically performs a message routing function and can be characterized as having a large number of signaling link terminations, a high volume of messages and the need for a large routing data base. The STP does not directly perform call processing functions and does not require connection to trunks. It is an independent function which, though intimately related to efficient operation of the toll network, is not directly associated with a toll switching machine. The initial STP has been developed for application in 4A/ETS offices which are not candidates for conversion to CCIS switching office operation. In this configuration the STP is intended to utilize the spare SPC processor capacity for signal unit processing and the Peripheral Bus Computer (PBC) for traffic mea-

surements and data administration. This arrangement has facilitated early introduction of the CCIS network. It is expected that the current STP will evolve into higher capacity configurations as the volume of signaling and other feature traffic increases.

2.1 Hardware organization and maintenance

2.1.1 General

Because of its functional independence the current STP configuration is designed to have a minimum interface, both hardware and software, with the host 4A/ETS switching machine. There are no interconnections between the 4A crossbar frames and the STP equipment. No modification of 4A crossbar equipment is required and no recompilations of ETS data are needed. New hardware additions are principally in the SPC peripheral unit area.

The STP office does require modifications to the SPC processor to provide two important new features: the "dead start" package and the Insulated Gate Field Effect Transistor (IGFET) store capability. The "dead start" package is a combination of hardware-software features which enhance the capability of recovering quickly and efficiently from a "dead" system. This includes such improbable multiple trouble situations as mate store failures and mate failures in critical peripheral units as well as mate processor failures.

The second new feature gives the processor the capability of working with the new IGFET store. The IGFET store is more economical than the Piggyback Twistor stores originally used in the SPC system and is used to provide the additional memory needed for the STP program and office data.

Any PBC-equipped ETS switching machine with the above processor enhancements which was not planned to become a CCIS switching office was a candidate for selection as an STP. Offices in this category were those scheduled for early relief by installation of a new No. 4 ESS,¹⁰ or those not highly interconnected to other offices which were CCIS switching office candidates. Additional criteria for STP site selection were that the office should be served by at least four physically independent transmission routes and each STP should be geographically remote from its mate. These considerations assure the most reliable signaling network and minimize the risk of a single catastrophe affecting mate STPs. There were at least two ETS offices in each of the ten regions suitable for STP selection.

The message handling capacity of an STP operating in combination with a fully loaded ETS is equivalent to one million call attempts per hour. The load carried by a region will normally be split between the mate STPs. In the event of the total outage of an STP, all the traffic for the

region must be carried by the mate STP. To insure the availability of this redundant capacity, each STP is normally loaded to no more than half of its capacity. The capacity of the region is then the same as the office capacity, or one million attempts per hour.

As the CCIS network matures and additional new No. 4 ESS machines join the toll network, the 4A/ETS switching machines that are coresident with the STPs will be relieved of their switching functions by the larger capacity No. 4 ESS machines. This step allows the STP to change to a stand-alone configuration with a subsequent capacity increase. The stand-alone STP will consist of the SPC complex, minimal peripheral and ancillary equipment, CCIS terminal equipment and the PBC complex. The capacity of a stand-alone STP is estimated to be two million attempts per hour.

2.1.2 Functional units

Figure 3 is a functional block diagram of the basic STP system. For simplicity, duplication of buses and controllers is not shown. For an explanation of the SPC-ETS complex see Refs. 2 and 3. The installation of an STP in a 4A/ETS office, in addition to the SPC processor enhancements mentioned above, requires the growth of equipment identical to that already provided for the ETS feature as well as the addition of new units developed for CCIS. The growth units include a peripheral scanner frame, a signal distributor frame and a central pulse distributor relay applique unit. In addition, power plant capacity enhancements may be required. The new circuits are the CCIS terminal group frame, the VFL access frame, the VFL test frame and a new unit added to the alarm and display frame.⁵ The additions to the PBC complex and its functions are described in Ref. 6.

2.1.2.1 Peripheral Scanner (PSC). This frame is identical to the existing ETS peripheral scanners and will be maintained by the existing ETS software. The STP-PSC will contain only scan points associated with the STP equipment, primarily the terminal frame out-of-service and link security alert points.

2.1.2.2 Signal Distributor (SD). The STP-SD is identical to the existing SPC-SD except that it is equipped entirely with individually controlled general-purpose relay circuits, as opposed to a mixture of specialized piggyback twistor store maintenance circuits and general-purpose circuits. Maintenance for this frame is provided by existing ETS software. The SD frame provides the out-of-service distribute point for the terminal units as well as the VFL access frame. This method of distribution was adopted in order to reduce cost and installation effort.

2.1.2.3 Central Pulse Distributor Applique circuit (CPDA). The STP-CPDA unit is identical to the existing ETS-CPDA and provides additional distribute points to augment those provided by the SD. These




Fig. 3—ETS/STP system block diagram.




Fig. 4—STP alarm and display panel.

points control the terminal access circuit out-of-service functions and the fifth peripheral scanner.

2.1.2.4 IGFET store. This frame provides the memory for the STP generic program and data. The IGFET store is designed for memory growth instead of the piggyback twistor memory used in current ETS store frames. ETS programs will remain in PET memory.

IGFET stores are duplicated for service reliability. One pair of IGFET frames should handle the requirements of the largest office. Each IGFET frame will mount up to six memory modules. The number of memory modules required in an office is dependent upon the number of CCIS terminals served by the STP office. One pair of modules will provide data memory for enough terminals to connect to nine other STP pairs and 27 CCIS switching offices. A second pair of modules will provide for connection to an additional 52 switching offices.

2.1.2.5 STP alarm and display panel. The STP requires a lamp panel to display the alarm the status information for the A, B, and C signaling links, terminal-modem units, and terminal access circuits (Fig. 4).

2.1.2.6 VFL access frame. The VFL access frame is a new frame designed specifically for the 4A STP to connect the terminal-modems to the signaling facilities, to provide VFL test access to the manual test position and to provide VFL test access to the test terminal unit (Fig. 2). In addition, the VFL access circuits contain adjustable transmission pads which provide the proper transmission levels at the modem and test points.

2.1.2.7 VFL test position. The VFL test position provides manual testing facilities for the CCIS voice frequency links. The position contains test equipment for making transmission impairment measurements of the type commonly used for testing private line voice-band data facilities. Each VFL appearance requires one four-wire jack and a lamp. Test access is provided under processor control via the VFL access frame by means of a TTY input message from the switching maintenance center.

2.1.2.8 CCIS terminal group frame. This complex contains the signaling terminal units and data modems for up to 16 signaling links, as

well as duplicated Terminal Access Circuits (TAC) for processor communication with each terminal unit via the peripheral unit bus. At present a maximum of 128 terminals (8 frames) can be equipped in an STP office, however, future enlargement to 256 terminals (16 frames) is anticipated. One terminal-modem unit in each STP office is reserved for VFL testing.

The terminal unit is a small, special-purpose stored program processor which maintains data communication over the signaling link. Synchronization, error detection, retransmission of signal units received in error, and acknowledgment of correctly received signal units are all handled by the terminal unit independent of the SPC processor. The terminal unit also provides multipriority buffering for incoming and outgoing signaling units.

The modem is the interface between the two-way serial digital data stream from the terminal and the two-way analog voice-band signal. One modem is physically and logically associated with each terminal. The modem has two VFL ports which, in the case of A links, are used to switch, under SPC processor control, between mate VFLs. One terminal-modem unit is associated with each signaling link. Backup capability for terminal-modem units is provided by the signaling network plan which requires duplication of all signaling links.

The fully duplicated TACs provide redundant independent access for each terminal unit via the ETS peripheral unit bus. The TAC has no autonomous functions and only one TAC is active at a given time. The SPC processor periodically polls the TAC to determine which terminal units, if any, contain waiting signal units.

2.1.3 Maintenance

The design of the signaling network and the strategy for reacting to signaling link failures was described in Section I. High reliability is an essential requirement of the signaling network. Redundancy is provided at several functional levels to assure uninterrupted service in the event of any single failure as well as a great many multiple trouble conditions. Basic to the signaling link recovery strategy is the assumption that the STP hardware is significantly more reliable than the VFLs. Highly reliable components have been employed in the terminal group and VFL access frames to minimize the component failure rate. The maintenance requirements for the STP are identical to those of the 4A/ETS switching office.⁹

The maintenance programs and procedures of the existing SPC complex are consistent with these requirements. As mentioned previously the ETS growth units added for the STP are maintained by existing ETS programs.

2.1.3.1 Detection. Fault detection in the terminal group hardware is accomplished by a variety of complementary and overlapping mechanisms. The TAC and terminal hardware contain extensive self-checking circuitry. The terminal software includes a self-test exercise program which runs continuously, interleaved with signal unit processing. Routine exercises are run on the TAC and terminal automatically on a daily basis to detect nonservice affecting faults. Faults in the modem and VFL access circuit are typically detected by excessive signal unit error rate or complete loss of the carrier signal.

2.1.3.2 Fault recovery. Faults detected in a TAC or terminal unit will cause a processor interrupt the next time the unit is accessed by the processor. This causes the terminal fault recognition program to be entered. This program determines which unit is faulty and then reconfigures the system with the faulty unit isolated. In the case of a TAC trouble, activity is switched to the mate TAC. In the case of a terminal trouble, a changeover request is made to the link security program. This causes the system to be reconfigured to use the mate signaling link as explained in Section 2.2. Fault recognition then schedules diagnostics on the faulty unit.

Faults which cause an excessive signaling link error rate lasting longer than three minutes will cause the link security program to request diagnostics on the suspect terminal-modem unit. If the subsequent diagnostics at both ends of the link find no trouble, a VFL failure is assumed.

2.1.3.3 Diagnostics and repair. Diagnostics are run on an interleaved basis with call processing. Their function is to isolate a fault to a replaceable component (circuit pack) and to verify circuit operation after the fault has been repaired (circuit pack replaced). The results of a diagnostic are printed on the ETS maintenance teletypewriter in the form of a trouble number. A Trouble Locating Manual (TLM) is provided for each unit type. It associates the trouble number with one or more suspected faulty circuit packs.

Diagnostic programs and TLMs are provided for the TAC and terminal. The terminal diagnostic includes a complete test of the modem and interface to the VFL access circuit. As such it represents the first completely automated test of data transmission terminal equipment incorporated into standard switching machine maintenance procedures. Special control and monitor features were incorporated into the modem design to make this possible.

Diagnostics for the TAC and terminal are written in a high level Diagnostic Language (DIAL).⁷ DIAL provides an efficient programming technique and a well structured, highly readable listing which facilitates the manual analysis of diagnostic test results if necessary. The TLMs, in addition to circuit pack information, also contain descriptions of faults

related to the interconnections between circuit packs. This information is very effective in locating wiring defects (shorts, opens) which sometimes occur on newly installed frames.

2.1.3.4 Maintenance terminal. As mentioned previously one terminal-modem unit in each STP office is reserved for VFL testing purposes. This terminal can be connected to any VFL via the VFL access circuit under control of the SPC processor. The main function of the maintenance terminal is testing of the reserve VFLs on A links. An automated processor-controlled test of the reserve VFLs is made on a scheduled routine basis to assure their availability when needed. This is a loop-back type test, initiated by the STP with a passive loop-back connected at the switching office. The maintenance terminal is also used to automatically perform the same type of test subsequent to a working mode VFL failure as a verification test.

2.2 Software organization

2.2.1 General

The STP program package is coresident with, but functionally independent from, the Electronic Translator System (ETS) programs. Both program packages are application related entities that operate within the framework of the basic Stored Program Control⁴ (SPC) programs. The SPC programs are composed of the executive control, maintenance control, interrupt, and input/output programs, and those programs required for the maintenance of the equipment providing the hardware core of the system such as processors, stores, central pulse distributors, and master scanners. These basic SPC programs are common to the Traffic Service Position System No. 1,⁸ the 4A/ETS, and the 4A/ETS/STP systems. In addition, they will continue to support the STP function in a stand-alone mode as the 4A machines at the STP sites are retired.

The STP program package provides for signaling message routing, signaling link security, and the administrative procedures to accomplish recent change and to collect traffic and plant counts. Maintenance software for the Terminal Access Circuit (TAC), the terminal, and Voice Frequency Link frame provide the necessary fault recognition and recovery mechanisms. Signaling link fault isolation and repair is facilitated by manually initiated software procedures, in addition to automatically administered visual and audible alarms. Audits and outage recovery strategies provide for software maintenance and complete the total maintenance picture.

The data base provided for the STP function is logically independent of the ETS data base and is generated and maintained separately. This data consists of routing instructions and signaling link information. The routing data consists of a table for each signaling link complement with

one word for each band containing the corresponding outgoing link and band number. Signaling link information contains link type, equipment arrangements and the identification of alternate routes.

When possible, development effort was reduced by providing common programs at both the 4A/ETS/CCIS Switching Office (SO) and the STP to serve common requirements. In addition to the SPC and ETS programs, the STP program occupies 18,000 words of memory of which 10,000 words are common with the 4A/ETS/CCIS switching offices.

2.2.2 Message routing

The primary function of the STP, signal routing, is performed by the message handler program which has been designed to optimize real-time efficiency. The fundamental process is to take an incoming message from one signaling link, change its band number, and transmit it on another signaling link as an outgoing message. The band translation and routing is accomplished by using the band translation table (Fig. 5) provided for each incoming signaling link. Using the band number from the incoming message as an index into the translation table, the outgoing band number and outgoing signaling link number are fetched from memory. The message handler replaces the incoming band number in the first signaling unit of the message with the outgoing band number, and places the message in an outgoing queue of the appropriate terminal. The outgoing link number from the translation table is the preferred route in the fixed routing scheme of the signaling network. Signaling link load balance on any given pair of B links to another region is done by assigning odd band numbers to one link and even band numbers to its mate.

Alternate routing on outgoing signaling links, if required because of an out-of-service condition of the preferred link, is achieved with real-time economy by retrieving an index stored in scratch memory for the outgoing link. The index effects a transfer to a routine which routes the message around a failed link or STP to its final destination. This index is called the transmit switch index, and is kept current by the link security programs so that the message handler may quickly make a disposition of any outgoing message.

In addition to band translation and alternate routing, the message handler is equipped to recognize signal units that initiate processing at the STP. As examples of this, a certain class of non-call-related messages causes the STP to broadcast Dynamic Overload Control (DOC) messages to prescribed switching offices on certain bands while other messages inform the STP of the suitability of a given signaling link to carry traffic. Most non-call-related messages, as well as regular telephone messages, are processed to completion once the incoming message is unloaded from a terminal. Lengthy sequences initiated by certain non-call-related messages are scheduled to be run as background jobs by the message




Fig. 5—STP message routing.

handler. The message handler itself is scheduled to execute as two separate tasks, a timed interrupt level job which runs every 10 milliseconds and is responsible for high-priority telephone traffic and a base level job which handles low priority telephone messages and noncall related messages.

2.2.2.1 High-priority message handler. Every 10 milliseconds the high-priority message handler is entered to look for high-priority traffic, currently consisting of answer and changeover messages each of which consists of one signal unit. The high-priority signal present indicators of each in-service Terminal Group (TG) are read looking for set bits which flag a terminal as having incoming traffic in its high-priority receive buffer. The message handler retrieves each signal unit, translates the incoming band, and routes answer messages to the proper outgoing terminal. Changeover events are queued for base level processing. The message handler continues to unload and process the messages from the active terminals associated with the current TG until all of the high-

priority buffers are empty at which time the high-priority signal present indicators return to zero. The message handler then moves on to the next TG. After all TGs have been served, the system returns to base level processing.

2.2.2.2 Low-priority message handler. The low-priority message handler which is entered on base level performs functions similar to the high-priority job with the added tasks of Multiunit Message (MUM) processing and filtering a wider class of message types which may require STP actions beyond routing. Upon entry, the message handler reads the low-priority signal present indicators of each TG looking for traffic queued in any of the low-priority input buffers of the 16 terminals. For every flag set the message handler retrieves the first signal unit of the incoming message and uses the first seven bits (the heading and signal information fields) as an index into a task dispenser which transfers control to a specific routine to deal with that message (see Fig. 5). A small percentage of the incoming traffic is trapped and processed at the STP, while the bulk of the messages—Initial Address Messages (IAM) and call-related Lone Signal Units (LSU)—are translated and routed. All LSUs that are routed through the STP are given the same treatment as described above for priority messages. Routed MUM traffic requires timed interrupt protected sequences for the transfer of the message from the incoming terminal to the outgoing terminal to prevent message inter-write. That is, once the outgoing link has been determined using the Initial Signal Unit (ISU) of the incoming MUM, each Subsequent Signal Unit (SSU) is retrieved from the incoming terminal and loaded into the outgoing terminal; the protected sequences prevent any timed interrupt generated low-priority traffic from being loaded prior to completion of the transfer of the last SSU. Incomplete MUMs are currently routed normally by the STP while any block of SSUs not preceded by an ISU is flushed from the network.

Incoming C-link traffic is always processed as a low-priority miscellaneous MUM since the message has been prefixed with a header signal unit which contains the outgoing signaling link number. This outgoing signaling link number is a result of translation performed at the mate STP which was unable to route the message due to the present network configuration. After stripping this header, a unique task dispenser is executed to process that particular message while implicitly recognizing that the message has been received from, and preprocessed by, the mate STP.

Each TG in turn has all low-priority traffic emptied from each of its 16 terminals which is then followed by an interject break to allow time-critical jobs to execute. After all TGs have been processed, the message handler relinquishes control to the next base level job.

2.2.3 Signaling link security

To ensure continuity of service, a large amount of redundancy has been designed into the CCIS signaling network. The signaling link security package at the STP has the responsibility of administering those links which connect it to other nodes in the network—this is the signaling link control function; and of maintaining and disseminating the status of signaling links between distant nodes—this is the signaling network management function. In addition to the obvious need to alternate route for failed links at the STP, the status of distant links is required on a band basis to ensure and optimize routing in the distant region.

Link security is responsible for monitoring the integrity of a working link, for initiating automatic restoral procedures for links subjected to terminal and facility failures, and for facilitating repair procedures on faulty links. Input stimuli to link security consist of signaling link messages, internal fault recognition, scan point notification, and manual inputs such as TTY messages and frame key actions. Broadly speaking, the resultant outputs from link security in response to a given stimulus may result in (i) changing the routing for some or all of the bands associated with a link, (ii) changing the link status, (iii) passing status information to distant nodes (switching offices, the mate STP, and/or distant STPs). The signaling link control function is outlined below and is followed by a review of the signaling network management function.

2.2.3.1 Signaling link control. The health of the signaling link is constantly monitored by the CCIS terminal. For example, the terminal is watchful for internal logic problems, an excessive number of signal units received in error, loss of block synchronization, and other abnormal conditions which may indicate a near-end signaling link problem. Diverse methods are used by the terminal to communicate the near-end problem to the processor. For one class of problems a special message, the processor notification signal, can be generated by the terminal and placed in a receive buffer to inform the processor of excessive signal unit errors, loss of block synchronization, or the buffer-full condition which arises when the last empty buffer slot is used for either a received signal unit or a signal unit loaded for transmission by the processor. Various control register or special data memory bits are read by the processor to identify the particular problem. In another class, the buffer overflow condition, which occurs during a buffer full state when the terminal attempts to put a received signal unit on one of the received link lists and finds there is no buffer space available, is reported to the processor by setting the Central Control Alert (CCA) miscellaneous scan point in conjunction with flags bits in the terminal control register. Finally, that rare class of severe terminal logic problems detected by regular terminal exercise routines forces the terminal to halt and thus deny the all-seems-well response to

the processor the next time the terminal is accessed. In this case, the peripheral fault recognition routines are entered.

In addition to monitoring for near-end detected difficulties, the terminal screens the incoming message stream for changeover signal units (COV) which indicate that the far-end has recognized a VFL problem and is in the process of switching traffic to an alternate link. Single COV units are discarded by the terminal; if two COVs are received within 256 signal unit intervals the second COV is passed to the processor in the high-priority receive buffer.

The normal changeover sequence is entered when a problem results in a changeover request and the mate signaling link is operational. Traffic destined to the troubled link is rerouted to the mate. All telephone and management messages waiting in the failed terminal transmit queues and all signal units transmitted but awaiting acknowledgment from the far-end terminal are moved to the mate link's transmit queues. Those messages waiting in the receive buffers of the failed link are unloaded normally by the message handler program. Link recovery procedures then advance the terminal through various states as resynchronization and prove-in are achieved. If the overall recovery timer expires, terminal diagnostics are requested and the craft personnel is notified.

The emergency restart condition arises when a changeover is requested on a signaling link and all other alternate signaling paths are not operational. Under these circumstances all signal units queued on the failed link are dumped while abbreviated recovery procedures are initiated. Emergency recovery prove-in thresholds are lower than those used during normal recovery in an effort to quickly regain signaling ability. During the emergency recovery interval, the affected DDD trunks are automatically inhibited from accepting new call originations.

Signaling link security at the STP is responsive to manual actions that affect signaling link states. Removal from service is a prerequisite for further maintenance activity such as manually requested terminal diagnostics. The manual changeover request can be initiated at either end of a signaling link. A request directly to the processor by the craftsman invokes the near-end procedure. Requests from the far-end are received over the signaling link as a Manual Changeover request message (MCO). Prior to sending the MCO on the link to be removed from service, the requesting office first inspects the status of the mate signaling link to ensure that loss of signaling will not result from a changeover. Before granting the request by sending the manual changeover acknowledgment, the receiving office goes through the same procedure. If both ends approve, traffic is removed from the link, but the terminals will maintain synchronization on the link. Should an emergency restart situation arise while a link is in the manual out-of-service state, link security can seize it and attempt emergency recovery procedures on it. Links manually

removed from service can be manually returned to service only by the requesting end or by both ends if a manual request was initiated at both ends. Once manually out of service, link security provides for a number of manually requested procedures to facilitate link fault isolation and repair. Examples are scheduling an automatic test of a VFL, connecting a VFL to the test position, or looping back a VFL at the far end.

2.2.3.2 Signaling network management. The signaling network management function is accomplished at the STP by administering status information on each band for each active B-link complement. Each of the 512 bands are represented by two bits in the band status table. One of the bits indicates that a Transfer Restrict message (TFR) has been received on that band from an STP in the distant region while the other bit indicates that a Transfer Prohibit message (TFP) has been received. The band status reflects the ability of the distant STP to route messages for that band and are used by the STP to make a disposition of outgoing messages on that band intended for one of the B-links of the pair. If there are no signaling restrictions on a band, the band is referred to as "allowed." Messages sent on restricted bands will require C-link routing in the distant region. Prohibited bands represent signaling paths that are blocked due to network failures. Upon receiving a band status message from a distant region the band status table for that incoming B-link pair may be updated and the appropriate switching office notified with a TFR or TFP. Band status information is generated by an STP based on link status changes of A-links and C-links at that STP and its mate STP. TFR and TFP restrictions for a given band are removed upon receiving a Transfer Allowed message (TFA). Additional procedures are provided for regenerating and updating the band status tables for links that have been out of service, as well as maintaining their integrity through audits.

Finally, link security with the cooperation of terminal software is able to minimize the effects of and recover from STP processor overloads and outages. Each terminal is able to recognize that the processor has not been servicing its low priority receive queue. The terminal then assumes that the processor is experiencing difficulties and autonomously informs the far-end node. The far-end processor, after an appropriate timing interval, effects a changeover on the signaling link to reroute traffic intended for the failed STP. As the processor at the failed STP proceeds through its recovery phases, the links are restored to service.

2.2.4 Maintenance features

The signaling link and network maintenance features discussed above are augmented by both hardware and software related processes to complete the maintenance package for the STP. Fault detection, fault recognition, and diagnostics for the TAC and terminal have been inte-

grated into the normal interrupt mechanisms provided by the SPC programs. These programs are common with those at the switching office as discussed in Ref. 9. In addition, the STP has been equipped with an Alarm and Display (A&D) status panel. Once alerted to the existence of a signaling link problem by the audible alarm, the A&D panel is available to the craft as a visual summary of the office troubles. The A&D panel classifies the troubles by unit type and severity. Should multiple troubles exist, both classifications are used in determining the sequence in which to perform the repair.

The display panel has two lamps (primary and secondary) providing three alarm states for each of the five categories (see Fig. 4). A steady secondary lamp indicates trouble in a single member of the unit type in question. A flashing secondary lamp means that two or more members of a unit type are in trouble, however, there is no loss in signaling capability. The minor audible alarm is sounded when this state is entered. The primary lamp is the most urgent warning. This state corresponds to multiple trouble conditions which cause service degradation. The major audible alarm is sounded when the primary alarm state is entered.

Software maintenance is provided by a series of audits which run on a scheduled basis, as time filler jobs, or on demand only. These audits monitor the status information in various link state tables. Discrepancies are automatically corrected and reports are made to craft personnel. Some audits involve the exchange of information between the STP and other nodes.

2.2.5 Administration

The STP has been provided with a comprehensive set of administrative procedures to facilitate data base maintenance and to provide for the collection of signaling network performance data. These features can be categorized into recent change, hardware growth, and traffic and plant measurements. An STP recent change function has been provided to update the band translation and other tables in order to establish or remove the cross-office signaling connections and overload or signaling congestion controls. This function provides for making changes to existing data tables that have been assigned in the data load. These tables may be associated with signaling links either preplanned or in service. The recent change administration can be accomplished via either the SPC maintenance channel or the Peripheral Bus Computer (PBC) port using similarly formatted messages which are oriented more toward the craftsman and less toward the actual structure of the data table being changed.

Another set of STP programs are used to administer the data state tables associated with the CCIS units. The data state concept allows the

installation and growth of CCIS units with minimal impact on an operational system. There are data state tables associated with the terminal access circuit, the terminal circuit, and the VFL access circuit. By changing the data state of a unit, it may progress from the unequipped state through the diagnostic state to the in-service state. Each of the states makes the unit progressively more available to an operating STP system.

Finally mechanisms diffused throughout the STP software collect traffic and plant counts which aggregate in the PBC. Examples of traffic data counts are signal units transmitted and received, initial address messages transmitted and received, and near-end and far-end changeovers. Plant data counts include signal units in error per VFL and retransmission requests per VFL. This information stored, analyzed, and reported by the PBC to both traffic and plant personnel contributes to the proper performance of the signaling network on a short-term and long-term basis.

III. ACKNOWLEDGMENTS

The design, construction and operation of the CCIS signaling network represents the efforts of many people throughout the Bell System. Without their contribution the network described in this article would not be a reality. The authors wish to thank S. D. Coomer for his assistance in the preparation of the STP software section.

REFERENCES

1. C. A. Dahlbom and J. S. Ryan, "History and Description of a New Signaling System," B.S.T.J., this issue.
2. B. T. Fought and C. J. Funk, "An Electronic Translator System for Toll Switching—System Description," IEEE Trans. Commun. Technol., COM-18, June 1970, pp. 168–175.
3. L. M. Croxall and R. C. Snare, "Electronic Translator System for Toll Switching—Routing Strategy and Network Control," Proc. National Electron. Conf., 25, 1969, pp. 400–405.
4. G. R. Durney, et al., "Stored Program Control No. 1A," B.S.T.J., 49, No. 10 (December 1970), pp. 2445–2507.
5. B. Kaskey, et al., "Technology and Hardware," B.S.T.J., this issue.
6. R. C. Snare, "Peripheral Maintenance and Administration Support System for CCIS," B.S.T.J., this issue.
7. G. A. Raack et al., "Development Tools," B.S.T.J., this issue.
8. R. J. Jaeger, Jr., and A. E. Joel, Jr., "TSPS No. 1—System Organization and Objectives," B.S.T.J., 49, No. 10 (December 1970), pp. 2417–2443.
9. K. E. Crawford et al., "Toll Crossbar Application," B.S.T.J., this issue.
10. L. M. Croxall, "No. 4 ESS Application," B.S.T.J., this issue.