THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 58, No. 6, July-August 1979
Printed in U.S.A.

Traffic Service Position System No. 1:

Automated Coin Toll Service: Software

By R. AHMARI, J. C. HSU, R. L. POTTER, and S. C. REED

(Manuscript received December 28, 1978)

The Traffic Service Position System (TSPS) No. 1 operational software for Automated Coin Toll Service provides the logic which controls the handling of coin-originated calls served by the system. The partitioning of responsibilities between the Station Signaling and the Announcement Subsystem (SSAS) are illustrated. The maintenance philosophy, fault detection, diagnostics, and fault recovery aspects of the SSAS are described. The maintenance strategy is centered on a multilevel fault detection scheme in which faults are analyzed and classified according to their degree of seriousness in affecting the SSAS operation.

I. OPERATIONAL SOFTWARE

Automated Coin Toll Service (ACTS) is a feature of the Traffic Service Position System.[1] A Station Signaling and Announcement Subsystem (SSAS) has been added to TSPS to detect and process coin deposits and to construct announcements for coin sent-paid toll customers.

The SSAS contains a programmable controller (a microprocessor), which has a Read Only Memory (ROM) for program and a Random Access Memory (RAM) for transient data. It also has a set of Coin Detection and Announcement circuits (CDAS). These circuits detect coin deposit signals from coin stations and decode digital speech phrases to produce analog announcements. The SSAS also has an Announcement Source and Distributor (ASD) which contains an announcement memory with digitally encoded speech phrases (see Fig. 1). Commands to the SSAS are sent by the TSPS processor over the

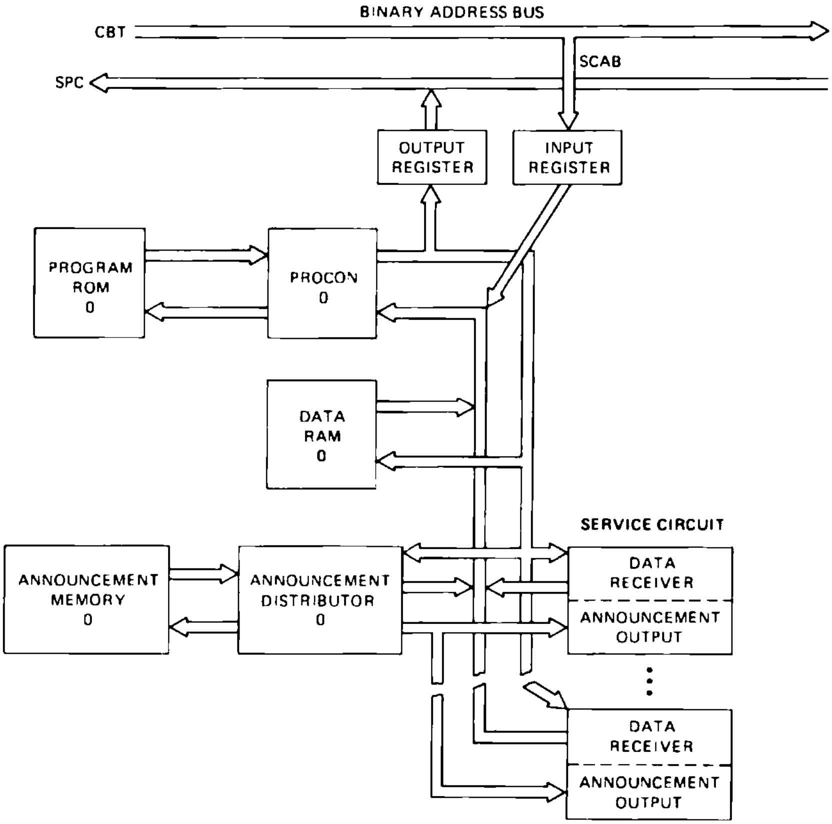

Fig. 1—SSAS configuration.

Peripheral Unit Address Bus (PUAB). Replies to the TSPS processor are sent over the Scan Answer Bus (SCAB).

1.2 Functional description of ACTS call processing

1.2.1 Initial call setup

When the TSPS call connections program* receives a report of an incoming trunk seizure† from the supervisory scan program, it establishes the required network connections for called digit and calling digit reception. (The latter is for Automatic Number Identification [ANI] offices.)

The ANI digit analysis program assumes control until reception of the calling party's number is completed.

* This program runs on the TSPS main processor.

† The description that follows deals only with calls that originate on trunks on the base network. Calls that originate on the network of an RTA are also handled by the operational software but not described below.

When the calling party identification has been received, control is returned again to the call connection program. At this time, a general analysis is performed on the information obtained, and the call is marked as 0+, 1+, etc. The 1+ coin-originated calls are candidates for automated treatment. (Coin customers expecting to make deposits to pay for a station-to-station call will dial the call with a "1" prefix or no prefix.)

1.2.1.1 Initial ACTS processing. The next step in processing the call is to determine whether the call can be automated. This is done by a program in the main TSPS processor.

Conditions for Automation. 1+ calls that satisfy the following criteria are candidates for automation during the initial contact on the call.

(*i*) ACTS-Converted Trunk Group. The call must be on an ACTS-converted trunk group. Certain modifications are needed in the coin station to generate dual-frequency coin deposit tones which the Coin Detection and Announcement circuits can recognize. All coin stations served by a trunk group must be modified before any calls on that trunk group can be automated.

(*ii*) Machine Ratable. The call must be machine-ratable; i.e., TSPS must receive or have in office data sufficient rating information to calculate the charges due on a call.

(*iii*) Not a Postpay Coin Originating Station. The call must not be from a postpay station. Coins deposited at postpay coin stations cannot be returned. An operator must verify that the correct party or station has been reached before the customer makes any deposit.

(*iv*) Not a Large Charge Call. The call must have an initial charge less than a certain threshold. Coin station hoppers handle only limited numbers of coins (24 nickels, for example). Operating practices instruct operators to make partial collections for every two to three dollars deposited if the call has large charges. Large charge calls require multiple collections. Hence, an operator is required to verify called party answer before any coins are collected.

(*v*) Automatic Number Identification. The call must have successful Automatic Number Identification. If an ANI failure occurs or the call is Operator Number Identified (ONI), the call cannot be rated. An operator must key the calling number.

Coin deposit monitoring on calls that fail condition (*i*) is not automated but is handled by operators using current procedures. However, notification at the end of the initial period can be automated on all coin calls, even if dual-frequency oscillators are not installed in the coin stations.

Calls that fail conditions (*ii*), (*iii*), or (*iv*), as well as person-paid, coin-originated calls (which are dialed as 0+ calls), seize a position for the required operator assistance. If the trunk group is ACTS-converted, a Coin Detection and Announcement circuit is attached to assist the operator in counting deposits. Subsequent deposit monitoring for overtime can be fully automated unless the large charge threshold is exceeded.

Calls that fail only condition (*v*) seize a position for the purpose of acquiring the calling number. After the number is keyed in by the operator, further handling of the call is automated.

If the above conditions for automation are met, the call connections program seizes an idle Peripheral Order Buffer (POB) and transfers control to the network control program which loads the POB with orders to establish a connection between the calling customer and a Coin Detection and Announcement circuit. If a CDA is not available, the call connections program seizes an idle POB and transfers control to the network control program which loads the POB with orders to establish a connection to both a position and an outpulsing circuit. The call is subsequently handled as a non-ACTS coin call.

SSAS Processing of Initial Deposit Requests. Processing in the SSAS begins when the Programmable Controller (PROCON) receives an initial deposit request command from the TSPS processor. The command is read by PROCON from the SSAS input registers. The command message includes the number of the CDA handling the call, the initial duration of the call, and the amount to be requested from the customer (see Fig. 2). For purposes of illustration, assume the charge is $1.15 and the initial period is three minutes.

The command field of the message is used as an index into a transfer table. The PROCON transfers to the address retrieved from the table. That program records the call data in the SSAS data RAM, sets the initial state indication for the call, and sends an output command to initialize the CDA circuit.

The PROCON next initiates scanning of the CDA circuit for coin deposits. The PROCON will continue scanning the CDA circuit at least once every 250 ms for the duration of the initial deposit phase of the call. The PROCON also instructs the SSAS announcement control cir-

	19	17 13	12 8	7 0	
WORD 1		INITIAL PERIOD	COMMAND	CDA CIRCUIT NO.	IRA
			12 9	8 0	
WORD 2			CREDIT	CHARGES	IRB

Fig. 2—Typical SSAS input message from TSPS processor. Message is in format as sent from TSPS processor (two 20-bit words).

cuitry to send the first segment of the appropriate announcement phrase to the CDA circuit. It sends subsequent commands for successive portions of the announcement every 512 ms. The announcement used for the call being described is:

"One dollar and fifteen cents,* please." (2-second pause) "Please deposit one dollar and fifteen cents* for the first three minutes."

If the customer deposits during the announcement, the CDA circuit instantly inhibits the announcement. When the PROCON scans the CDA circuit and recognizes the deposit, it ceases sending further announcements, adds the value of the coin deposit to the previous amount deposited, and compares the total with the amount due. Assuming a sufficient deposit has not been made, the PROCON begins timing. If the customer fails to deposit within *five or six seconds*, a prompting announcement is provided indicating the amount still to be deposited. For example, if the customer deposits three quarters in the above example and then stops, the prompt is:

"Please deposit forty cents more."

If the initial deposit announcement completes without a deposit, timing will begin at that point. If the customer has made no deposits and a time-out occurs, the announcement wording for the prompt is:

"Please deposit one dollar and fifteen cents."

Each deposit made causes the PROCON to reset its software intercoin timing register for the call.

The PROCON will report the final results of the initial deposit request to the TSPS processor by loading a message into the SSAS output FIFO buffer (see formats in Fig. 3). There are basically three situations possible.

(*i*) If the customer has failed to deposit within five to six seconds after a prompt, the PROCON sends a reply (Fig. 3, format B) to the TSPS processor indicating this fact so an operator can be connected to provide assistance.

(*ii*) If the customer deposits the exact amount requested, the PROCON sends reply format A which contains the amount deposited.

(*iii*) If the customer overdeposits by using too large a denomination coin, the PROCON sends reply format A. (The TSPS processor

* The announcement format varies slightly, depending on whether the charge involves a dollar amount, a cents amount, or both.

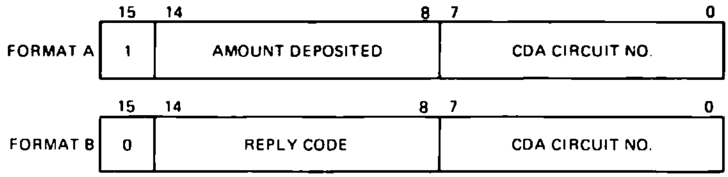

Fig. 3—Formats of SSAS output FIFO buffer.

recognizes the overdeposit and records a credit towards overtime.)

The PROCON acknowledges the latter two cases above by initiating announcements. If the exact amount is deposited, the announcement is:

"Thank you."

If an overdeposit has been made (for example, $1.25 on the $1.15 call described above), the announcement is:

"Thank you, you have ten cents credit towards overtime."

The PROCON continues to scan the CDA for further deposits until the acknowledging announcement is completed and sends a final deposit report to the TSPS processor at the end of the announcement. (This ensures credit for a belated deposit.) Finally, the PROCON places the CDA and the associated call memory in the SSAS data RAM in the idle state. The PROCON performs no further action on the CDA until a new command is received from the TSPS processor.

At any time during the processing described above, the TSPS processor can send a special command which causes the PROCON to idle the CDA circuit and associated call processing memory. (One example of this happening is if the customer hangs up.)

Successful Initial Seizure. Upon receipt of the reply from SSAS indicating an exact deposit or overdeposit, the call is processed to completion. The call connections program seizes an idle POB and transfers control to the network control program which loads the POB with orders to connect an outpulsing circuit. If an outpulsing circuit is not available, the call connections program queues until the circuit is available. The call connections program loads the orders to perform the appropriate relay operations required to complete the connection and activates the POB. Upon successful POB completion, control is returned to the call connections program where the POB is idled. The outpulser loading routine is called next. It loads the digits to be outpulsed in an outpulsing register and activates sender-attached scanning for the receipt of a sender-attached signal from the toll office. Outpulsing of the called number to that toll office proceeds as described in Ref. 2, page 2658 ff.

Receipt of called party answer and call timing also follow the description in Ref. 2.

Operator Assistance on Automated Initial Seizure. If the customer fails to deposit in response to prompting by the SSAS or flashes the switchhook* to acquire operator assistance, the customer will be connected to an operator. This section discusses the display presented to the operator, the operator actions, and potential race conditions.

If the call must queue for a position, it is given "recall priority" to minimize customer delay. The Coin Detection and Announcement circuit is connected to the calling customer during the queuing interval. If the call is being sent to a position because of a time-out and the customer subsequently satisfies or exceeds the charges while queuing for a position, the SSAS informs the TSPS processor, the call is removed from the position queue, and outpulsing is initiated without operator assistance. When the call reaches the position, the following keys and lamps on the operator console (see Fig. 4) are lit steadily.

The loop access key (ACS)
The appropriate supervision lamps (CLD), (CLG)
The station coin lamp (STA)
The AMA station-paid key (PAID)
The release forward key (FWD)

The lighting of the AMA station-paid key and the release forward key indicates to the operator that an ACTS time-out or a customer switchhook flash has occurred during an ACTS initial contact.

In addition to the above, an ACTS underdeposit display is given to the operator (see Fig. 5). The "charge-minute" designation strip is lit steadily and the numeric field contains, from left to right, up to three digits for charges due, one digit for the initial period and up to three digits for the amount due.

While the call is attached to a position, the operator assists the customer in making a full deposit. The operator must announce the amount due (the right-hand quantity) in the numeric field. The CDA circuit monitors and counts the coin deposits. The numeric display is not automatically updated with each coin deposit. However, the operator can update the display by using the charge and minutes (CHG-

* For a time-out, the SSAS indicates the request for an operator by a message. For a flash, the TSPS processor initiates the request.

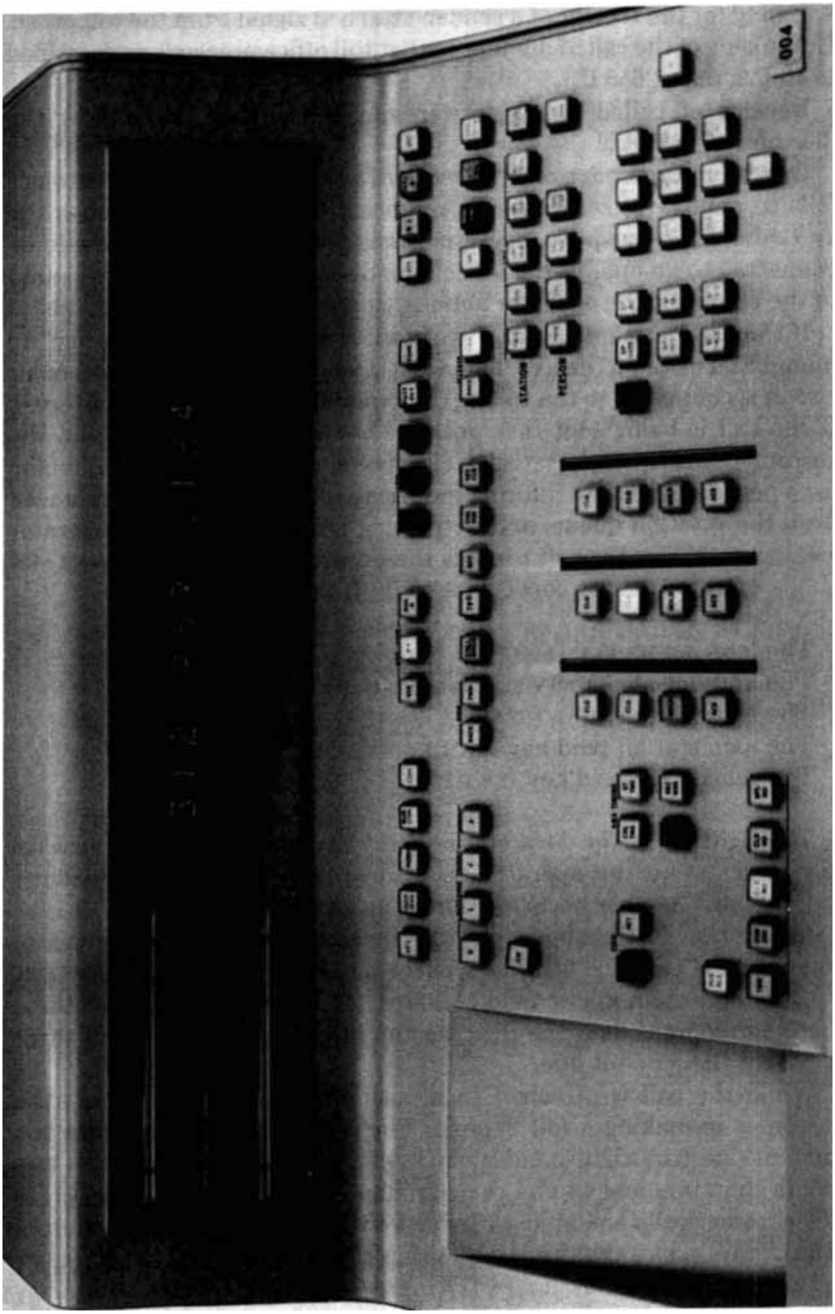

Fig. 4—100B traffic service position key shelf.

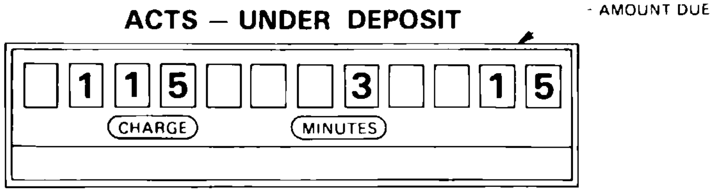

Fig. 5—TSPS No. 1 console numeric display.

MIN) key. (The SSAS will be interrogated for the current amount deposited and the appropriate display presented.)

When the SSAS recognizes that the deposit request is satisfied, the TSPS processor is informed. The forward number is outpulsed and the correct display is presented to the operator (see Fig. 6).

The operator receives several indications that the deposit is satisfied: outpulsing occurs, the amount due display changes to zero, the release forward key darkens, and the start key lights during outpulsing.

After acknowledging the deposit, the operator depresses the start timing* and position release keys.

At position release, the TSPS processor requests the amount detected by the CDA circuit. Upon receipt of the report, the CDA circuit is disconnected from the call. This final report accounts for all money deposited (while the CDA was connected to the call) so no late deposits are missed.

If an overdeposit occurs, the numeric display is changed to an overdeposit display (see Fig. 7). This display has a flashing "charge-minutes" designation strip and a numeric display of zero in the charge field, blanks in the minutes field, and two digits having the overdeposit. The operator informs the customer that credit will be given towards subsequent overtime charges on the call, depress start timing, and position release. On subsequent deposit requests, the overdeposit is automatically subtracted from the charges due.

1.2.1.2 Coin notification. At the end of the initial period, coin sent-paid customers are notified. With ACTS, this no longer requires an operator. The description below applies to coin sent-paid calls from both ACTS and non-ACTS converted stations.

TSPS Processing of Coin Notification. When a call reaches the talking state, it is under the control of the disconnect program. The call is placed on a timing list when the called answer has been established. A time-out occurs 7 seconds prior to the end of the initial period. Subsequent to this time-out, a coin collect sequence† occurs

* Timing on the call is initiated by the TSPS processor only if the called customer answers.

† The coin collect sequence does not occur on postpay coin stations.

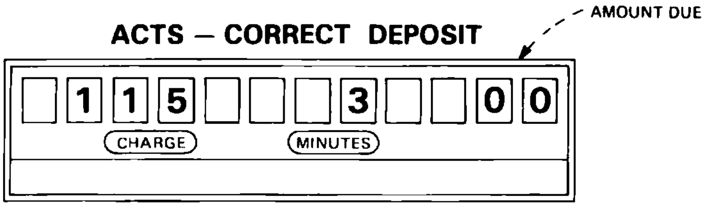

Fig. 6—TSPS No. 1 console numeric display.

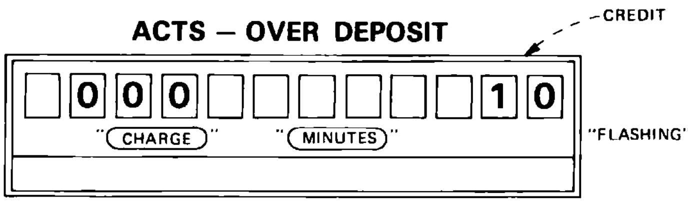

Fig. 7—TSPS No. 1 console numeric display.

and the call is returned to the timing list for the remainder of the initial period.

At the end of the initial period, the disconnect program seizes an idle Peripheral Order Buffer (POB) and transfers control to the network control program which loads the POB with orders to establish a connection between the calling customer and a Coin Detection and Announcement circuit.

SSAS Processing of Coin Notification. When the PROCON receives an input message requesting an end-of-initial-period notification on a coin call, the data accompanying this command are the CDA number and the time of the initial period. In response to this command, the PROCON directs the SSAS announcement control circuitry to give the following announcement:

"Three minutes has ended, please signal when through."

(A three-minute initial period is used for this illustration.) No coin scanning is performed, as no deposits are expected. After the announcement, the CDA hardware and software are idled and an announcement complete reply code (Fig. 3, format B) is sent to the TSPS processor.

1.2.1.3 Charges due seizures. The ACTS procedures for fully automating overtime charges due seizures on coin-paid calls are presented in this section.

Conditions for Automation. The SSAS can fully automate overtime charges due seizures on coin-paid calls on an ACTS-converted trunk group even if the initial contact requires an operator. Since the call is already rated, no rating or calling number identification (ANI) restrictions apply to overtime seizures. Furthermore, charges due seizures for postpay calls are handled in the same manner as calls from other coin stations. However, there is one additional restriction on charge due seizures. All previous deposit request seizures (either initial contact or charges due) must have been successfully monitored by a CDA circuit. Successful monitoring implies that CDA circuit blocking did not occur and that the CDA circuit did not malfunction. This condition ensures that overdeposit credits are not lost. (If a seizure is not monitored successfully and an overdeposit occurs, the credit is not recorded.) If successful monitoring does not occur, the call is processed by an operator using current practices. In summary, the conditions for automating overtime charge due seizures are:

(*i*) The calling station is on an ACTS-converted trunk group.
(*ii*) The charges do not exceed the large charge threshold* (see Section 1.2.1.1).
(*iii*) All previous seizures were successfully monitored.
(*iv*) CDA must be available.

If a call continues for a certain number of overtime intervals (usually 10) and all the conditions for automation are satisfied, the SSAS automates the changes in seizure. Additional intermediate deposits are requested at successive specified overtime intervals until the call terminates. Then the SSAS automates the final charges due seizure at the end of the call. The collection sequences for end-of-call deposits and intermediate deposits are essentially the same. Furthermore, these collection sequences are very similar to those used for the initial contact on station (1+) calls (Section 1.2.1).

With only a momentary interruption to the conversation path, TSPS connects an idle CDA circuit to the call. TSPS informs the SSAS of the amount due, the number of minutes which have elapsed, the CDA circuit being used, and that it is a charge due seizure.

SSAS Processing of Overtime Deposits. SSAS deals with the two overtime situations (end of call and intermediate overtime deposits) in a similar fashion. After receiving the command from the TSPS processor, the PROCON processes the call in a fashion similar to the initial deposit case and the reports returned to TSPS are also similar. However, the TSPS sends the amount of credit from previous deposits (if any) and the time sent is the elapsed overtime talked (in minutes), not the initial period. The announcement wording is also different. Assume that the call has progressed for five 25-cent overtime periods of two

* The office data provided for the initial-period large-charge threshold is also used to specify the overtime large-charge threshold.

minutes duration each and that the customer had no credit. The announcement would be as follows:

(Alerting tone) "Please deposit one dollar and twenty-five cents." (0.5-second pause) "One dollar and twenty-five cents for the past ten minutes."

If the customer had a ten-cent credit, for example, the announcement is altered.

(Alerting tone) "One dollar and fifteen cents please." (2-second pause) "You have ten cents credit. Please deposit one dollar and fifteen cents more for the past ten minutes."

In the intermediate overtime deposit case, both customers are on the call and intend to continue talking. They both can hear the announcement and talk during the deposit interval. In the end-of-call overtime case, only the calling customer is on the call.

As with the initial deposit situation, if a customer stops depositing coins, a prompting announcement is given. Failure to respond to this announcement results in a report to TSPS requesting an operator.

Operator Assistance. If the call times out or the customer flashes the coin station switchhook during the charges due phase of a call, an operator is connected to the call. The following keys and lamps on the operator's console are lit steadily:

(*i*) The loop access key (ACS).
(*ii*) The appropriate supervision lamps (CLG, CLD).
(*iii*) The appropriate coin lamp (STA, 0+, or Dial 0).
(*iv*) The appropriate AMA key (station paid or person paid).
(*v*) The charges due (CHG-DUE) lamp.

These are the same keys and lamps lit by TSPS on charges due seizures prior to ACTS. In addition to these keys and lamps, the ACTS underdeposit display, previously described for initial contact, is lit on the operator's console. The appearance of the ACTS underdeposit display indicates two things to the operator. First, a CDA circuit is attached to the call to monitor the deposits. Second, if the customer overdeposited on a previous seizure, the charges displayed have been corrected by the amount of credit. (Note that, on non-ACTS charge due seizures, the operator must subtract any credit claimed by the customer from the charges displayed.)

The operator assists the customer until a full deposit is received. When the SSAS recognizes that the deposit request is satisfied, the correct deposit display is given. The operator thanks the customer and depresses position release. Upon releasing the position, TSPS requests the amount detected by the CDA circuit. After TSPS receives this information from the SSAS, the CDA circuit is disconnected.

If the customer overdeposits during an overtime collection with an operator attached, the overdeposit display is lit on the operator's console. If this is a charges due seizure for intermediate collections, the operator should inform the customer that credit will be given on subsequent overtime charges. If the overdeposit occurs at the end of the call, the operator should handle the disposition of the overdeposit in accordance with current operator practices.

If the calling customer hangs up during the ACTS charge due announcement sequence, or the calling customer leaves the phone off-hook, the operator should try to reach the calling customer, if necessary, by ringing back against off-hook. The operator may need to wait for the customer to return. If the operator is unable to obtain full payment for the call, the operator keys a walkaway trouble number. A traffic counter is pegged, and a special bit and the amount of shortage is recorded on the AMA tape.

Even after the decision is made to go to an operator, the CDA circuit still monitors coin deposits. If the deposit is satisfied before the Peripheral Order Buffer (POB) which connects the position is activated, the position seizure is aborted. The SSAS acknowledges the deposit, and the call is released.

If the position is seized and then the deposit is satisfied, the SSAS does not acknowledge the deposit. When the operator's numeric display is lit, it shows (or may quickly change to show) that the deposit is satisfied. The operator should acknowledge the deposit and release the call.

1.2.2 Non-coin features

The announcement without coin detection mode of operation is used to provide announcements for time and charges quotations and notification on calls other than coin sent paid. A different command is used in each case to select a different announcement.

1.2.3 SSAS as a coin detector

The PROCON is also programmed to provide a mode of operation in which it gives no announcements but does monitor for coin deposits. This is used on operator-handled ACTS calls. In this mode, the PROCON informs the TSPS processor when the required deposit or an overdeposit has been made. If the operator collects or returns coins while a call is in this state, TSPS informs the SSAS so the PROCON can appropriately update the amount deposited and the amount due in its data RAM. If the operator requests that the amount deposited be displayed, the SSAS will transmit that amount to the TSPS processor in response to the appropriate command.

1.2.4 Coin station maintenance and administrative features

One unique aspect of ACTS operation is that inband signals (the coin deposit tones) are transmitted directly from the station through to the TSPS over the voice path. Two special features were developed that are related to coin tone signaling. First, a coin station test capability was developed which allows a craftsperson at the station to test that coin deposits for that station can be detected at the TSPS. Second, a precutover mode of operation was devised to help detect stations that have not been converted to be compatible with ACTS. These features are described below.

Coin station test call. After the craftsperson has performed the usual coin station tests, the TSPS is accessed by dialing a special test code. When TSPS receives this number, it connects the call to an idle CDA and sends a message to the SSAS with a command code indicating that this is a station test call. The remainder of the call is handled by the SSAS except that TSPS supervises the call. If the craftsperson goes on-hook, the TSPS aborts the call by sending a command to SSAS. Also, TSPS will return coins on request from the SSAS.

When the PROCON receives the station test call command, it initializes the CDA, begins coin scanning, and initiates the following announcement:

"Coin Test." (1-second pause) "Please deposit nickel."

If the craftsperson deposits a nickel, the announcement,

"Nickel"

is given, acknowledging the deposit. A dime and quarter deposit are also requested and acknowledged. Then, the craftsperson may make additional deposits which will be acknowledged if received correctly. If no coin is deposited for about six seconds, the deposit request is repeated. If the wrong coin is deposited, a coin return is requested by the PROCON and the test is repeated. If repeated requests for deposits are not satisfied or a total of two minutes elapses on the call, the PROCON issues instructions for the announcement.

"Test has ended."

and reports the time-out to the TSPS, which gives a coin return and disconnects the call.

1.3 PROCON program design

The PROCON has relatively high processing capacity. It has a relatively powerful instruction set which is largely at the assembly language level rather than the microprogramming level, and it manipulates a 16-bit word. It uses a paged program memory and can execute instructions within a page in arbitrary, fixed order. (The next interpage address (displacement) is specified in each 24-bit instruction.) Because of this characteristic, care must be exercised in relocating code to avoid improperly crossing a page boundary. (Such errors are detected during compilation.)

The PROCON's 16K program memory limit requires that the programs be very compact. While an attempt was made to use modern, structured, design and documentation techniques, the limited memory size caused compromises in some areas. (The SSAS application requires a far larger program than most PROCON applications.)

Multiprogramming was achieved by storing all call- (process) related information in the data RAM. Each program retrieves all data from that memory when it begins processing a new call. Each program is a pure process stored in ROM (i.e., the program contains no variable data and cannot alter itself).

The PROCON has very extensive self-checking features. In addition to a variety of parity checks (including next address parity), the arithmetic and logic circuits are duplicated and checked by matching.

1.3.1 PROCON program administration and design standards

Several designers developed separate programs which were integrated into one large program. Loading was accomplished by using a linkage editor and a common file of symbols and macros. All symbols (except labels) and all macros were defined in the common file. No numerical references were permitted, and input/output control designation symbols were keyed to the circuit drawing names for the same leads. Thorough prologue comments were required for each routine. Standard labels and standard symbolic designations and usage for all PROCON internal registers were agreed upon by all programmers.

1.3.2 PROCON active monitor

The basic program loop for the PROCON in the active SSAS (i.e., the SSAS which is processing calls) is called the active monitor. The PROCON executes this program while looking for tasks to perform. In this loop, the PROCON first makes some maintenance checks and then reads a clock to determine if a new 512-ms period has begun. The 512-ms interval is called a base period. Its length is keyed to the length of an announcement speech segment. If a new base period has begun, a special flag is marked in the system status register.* The PROCON then scans for a message from the TSPS processor. If such a message has arrived, it records the data in the appropriate RAM area.

The PROCON then begins a loop in which each CDA is processed. Each CDA has a dedicated eight-word call register associated with it. (Figure 8 is a typical CDA register layout.) First, the PROCON reads the address of the program which should process the call. This address, which also defines the state of the call, is called a Progress Mark (PM). If the PM is zero, the CDA is idle or unequipped, so no action is taken. A typical PM program is the program which scans a CDA for coins. The PM is stored as a page number and a displacement within the page. The active monitor transfers to the PM routine. If the PM routine determines that the call requires timing, it checks the new base period flag. If that flag is set, the required timing counter is updated.

When the PM program returns to the active monitor, the monitor checks the new base period flag. If that flag is set, all timing counters for the CDA are updated as required and a new announcement segment is selected. (The announcement processing is described in Section 1.3.4.) After the announcement program is completed, the monitor processes the next CDA. When all CDAs have been processed, the monitor clears the new base period bit, checks for various maintenance tasks, and then begins the main loop again.

1.3.3 Announcement store layout

Each announcement speech segment consists of 16,000 bits of digital data stored in 400 consecutive announcement store locations. Each location contains 40 bits of data. Decoded at a 31,250 bits/second rate, this corresponds to 512 ms of speech. For programming ease and circuit design convenience, all segments must begin at an address which is a multiple of 16. Most words used fit in one segment (for example, the numbers "one" to "ten"). However, several words or phrases require two segments (for example, the word "eleven" or the phrase "please deposit"). A few phrases require three segments (for example, the phrase "please signal when through").

Because the announcement circuitry bit-synchronizes segments, two or more segments can be blended together with no loss in quality. Therefore, words can be split between 512-ms segments to achieve the best cadence of speech.

Successive segments constituting 1-second and 1.5-second phrases need not occupy consecutive announcement store addresses. However, these segments generally must be used together. In contrast, suffix segments used for compound numbers may be combined with a variety

* One of the PROCON general registers was reserved for critical status indicators.

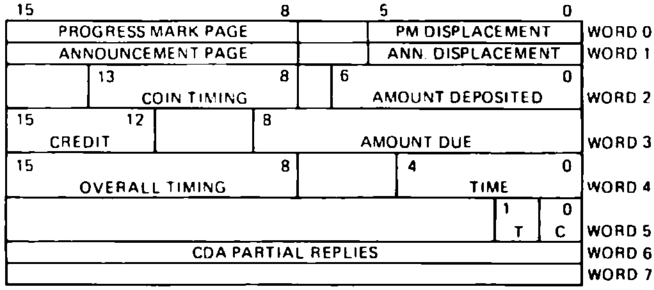

NOTES:

1. PM = PROGRESS MARK
2. AMOUNT DEPOSITED, AMOUNT DUE AND CREDIT IN NUMBER OF EQUIVALENT NICKELS
3. C–TYPE OF CDA
4. T TRUNK TYPE (2 WIRE = 0,4 WIRE = 1)

Fig. 8—CDA data RAM call register.

of other segments. For example, the suffix "-teen" can be combined with the syllables "thir-," "four-," etc. to form the numbers "thirteen," "fourteen," etc. with minimal use of memory. Similarly, the units place suffixes "-one," "-two," etc. can be combined with the numbers "twenty," "thirty," etc. to form the full required set of numbers (e.g., twenty-one, twenty-two, thirty-one, thirty-two).

Several words are duplicated in the announcement store so that both neutral and falling inflections are available. The falling inflections are used only at the ends of sentences, while the neutral inflections are used elsewhere.

In total, about 80 speech segments 512 ms in length are required for ACTS call handling. About 15 more are required for maintenance purposes.

1.3.4 Announcement table structures

Two major announcement tables are stored in the PROCON ROM. The first is simply a table of addresses. Reading this table using the speech segment number as an index provides the address of the first announcement store location containing the data for that speech segment.

The second table is more complex, as it controls the sequencing of announcement segments to form sentences. The logic for this process is embedded largely in this table rather than the program code. This makes debugging easier and also simplifies the process of adding new announcements. The table consists of program addresses for generalized routines dealing with each step of the announcement. The address of the current table entry is stored in the second word of the CDA register (see Figure 8). Each announcement routine must decide what speech segment is to be reproduced next and also update the table address in the CDA register. The major situations the various routines must deal with are demonstrated by the following example.

The structure of a portion of a typical announcement is represented schematically in Fig. 9. This is the end of the time-and-charges announcement used for noncoin calls. A typical example of this announcement follows.

> "The charges are one dollar and twenty cents plus tax for seven minutes."

The three major types of announcement routines are each involved in this announcement. One type of routine simply sequences speech segments in fixed order (for example, the words "plus tax for"). The second type of routine chooses a word or sequence of words ("one minute" versus "*M* minutes") based on the value of a parameter (whether *M* is one or greater than one). The third type of routine translates a parameter (*M*) into a speech segment (for example, the number "seven") or a series of segments (for example, "twenty-one"). This last program is a subroutine since the action to assemble the number is the same whether it represents a dollar amount, a cents amount, or a time period.

II. MAINTENANCE

2.1 Overview

2.1.1 TSPS maintenance requirements

In TSPS, as in other Bell System electronic switching systems, the maintenance strategy is based on duplicating vital hardware units[3] and providing hardware checking circuits for other units, using signals to indicate the successful execution of orders and using programs to test the state of the hardware, to detect faults, and to diagnose trouble.

TSPS provides fault recognition and diagnostic programs for all major circuit elements. The purpose of the fault recognition programs is to quickly detect faulty equipment units and, if necessary, remove them from service. Hardware checking circuits are used for trouble detection

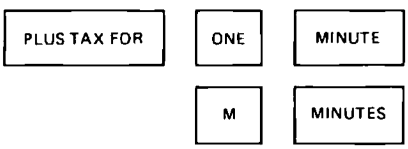

Fig. 9—Typical portion of an announcement phrase.

during operation. When a trouble detection circuit identifies a serious problem in the system, it notifies an interrupt circuit. The interrupt circuit immediately stops operational program processing and transfers control to a fault recognition program associated with the particular trouble indication. The functions of the fault recognition programs are to distinguish nonrepeating troubles (errors) from repeating troubles (faults), and, in the case of faults, to quickly establish an operational system configuration. This is generally done by switching out the faulty unit. When the fault recognition process is completed, operational program processing is resumed at the point of interruption.

Fault recognition is the highest priority function in a real-time system like TSPS. It is of the utmost importance to high system reliability that minimum time be taken away from the operational processing functions because of faulty circuits.

Once the faulty circuit has been taken out of service by the fault recognition program and the system brought back to an operational configuration, diagnostic programs are called in to test the out-of-service unit to isolate the trouble to a small number of circuit packs. A diagnostic program is normally broken into logic testing entities called subphases or phases so that a series of rigorous tests can be performed on the out-of-service circuit unit. Depending on the outcome (pass-fail) of each test, a trouble-locating number is printed on the maintenance teletypewriter. This trouble-locating number is used by the maintenance personnel to reference a trouble-locating manual to isolate the particular fault involved.

2.1.2 SSAS maintenance requirements and anticipated reliability

SSAS reliability and maintainability objectives are based on the fact that a significant part of TSPS traffic is served by the SSAS. The objectives are

(*i*) An average of about one SSAS outage in 10 years.
(*ii*) An average repair time of about 2 hours.

The maintenance plan for SSAS impacts on the reliability objective in three major areas. First, it is necessary to provide fault detection mechanisms (primarily hardware checks) so all failures can be detected rapidly. Second, provisions must be made to minimize the interruption of call processing once a failure is detected, with few or no calls lost. Third, sufficient circuitry must be provided to obtain diagnostic resolution to hold repair time to the 2-hour average as specified above.

The SSAS maintenance strategy relies on a multilevel fault detection scheme in which faults are classified according to their impact on the SSAS operation. A fault that does not affect normal operation of the

SSAS is considered "tolerable." An SSAS side with this type of fault may continue to perform its normal function adequately as long as it is required to do so. As an example, all single bit failures in the SSAS semiconductor announcement memory are tolerated, since they would have very little effect upon the quality of the speech. By using a maintenance strategy that permits an SSAS side to be used with certain "tolerable" faults, the SSAS reliability is further enhanced.

To verify that the SSAS design meets the reliability objectives, a reliability estimate was made by classifying and counting the boards on the SSAS side. The failure rate for each board type was calculated. The failure rate determines the Mean Time Between Failure (MTBF) of an unduplicated SSAS side. The Estimated Downtime (EDT) for the duplicated system was then calculated with the assumption that the mean time to repair for one side down and both sides down are identical. This is a valid assumption because the diagnosis of all critical circuits in SSAS can be done on each side independently. Plugging necessary parameters into the appropriate reliability model showed that the above reliability objectives are well satisfied.

2.2 Announcement Store Maintenance

2.2.1 SPC bus and store switch

The Stored Program Control No. 1A (SPC 1A) is the TSPS main processor.[4] The SPC memory is used for both programs and data. It is partitioned into a maximum of 24 duplicated stores. Each store is permanently assigned to one of the two SPC store buses. An SPC store can be equipped with a piggyback twistor (PBT) or semiconductor memory. Two sequential or complimentary name codes are assigned for an SPC store when equipped with semiconductor memory and one name code when equipped with PBT. The name code serves to identify the stores for addressing purposes. A semiconductor store consists of a semiconductor memory module and its associated semiconductor memory controller, shared by all modules in the same memory frame. The SPC store is the maintainable unit of SPC memory; i.e., it can be removed from or restored to service, diagnosed, and updated with the contents of its duplicate.

The SSAS announcement memory is used to store digitized announcements. An Announcement Store Frame (ASF) is permanently assigned to each side of each SSAS. The SSAS ASF consists of an announcement memory controller and up to six announcement memory modules. To simplify development effort and to take advantage of the memory loading facilities and diagnostic software of the SPC, the SPC semiconductor store frame with a second access port is used as the announcement store frame. This is accomplished by placing a switch on the announcement memory which allows the memory to be connected to either the SSAS side or the SPC store bus. Furthermore, to reserve the maximum number of memory name codes for the SPC store, all SPC memory access operations use odd parity name codes while all announcement memory access operations use even parity name codes.

Under the normal mode of operation, the switch is set to the SSAS side so that access to the announcement memory can be performed. In the SPC direct access mode, the switch is set to the SPC side so that loading from the Program Tape Unit (PTU) and diagnosing by the SPC can be accomplished. More discussion on loading and diagnostic operations is presented in a later section.

2.2.2 The choice of announcement memory

To handle all TSPS station-paid coin call announcements as well as time and charge quotations for noncoin calls, it was calculated in 1974 that announcement capacity equivalent to about 80 half-second words would be needed. An ultimate memory capacity of 200 to 400 half-second words was forecast at that time if vocabulary for future features as well as space for other TSPS announcements were considered.

A number of storage media were examined for the announcement memory. After some consideration, it was decided that with minor modifications the semiconductor memory used for the SPC 1A store would also be suitable for the SSAS announcement store. This choice of the announcement memory would provide flexibility in areas such as future vocabulary growth and vocabulary reload in the field from a magnetic tape. It was also estimated that most hardware and maintenance software designed for the SPC 1A store could also be applied to the announcement store with only a modest amount of effort.

2.2.3 Diagnostic implementation

The SPC store diagnostic programs were modified to be used for the SSAS announcement store diagnostic. Major objectives for the announcement store diagnostic implementation were specified as follows:

- (*i*) Circuits common to both designs to be tested by existing SPC semiconductor memory diagnostic programs.
- (*ii*) New announcement store circuits to be tested by the SSAS controller diagnostic.
- (*iii*) Same Trouble Locating Manual (TLM) to be used for SPC and SSAS announcement stores.
- (*iv*) Similar TTY input/output messages.

Several diagnostic characteristics are common for the SPC and announcement store diagnostic programs: the basic diagnostic unit is a controller and one memory module, seven diagnostic phases are in the program, the diagnostic terminates on early detection of errors, there is a special bus configuration for increased fault detection, and finally, there are automatic exercise tests for all 32K words. For these common areas, existing SPC store diagnostic programs were applied to the announcement store with only very modest modifications.

On the other hand, some diagnostic characteristics pertaining only to the SSAS announcement store must be added to the existing store diagnostic programs. First, the announcement store has to be switched onto the standby SPC store bus before it can be diagnosed. Second, since during normal operations, the PROCON and ASD circuits in the SSAS controller control the operation of the announcement store, they must be stopped during SPC diagnostic tests. Third, the SSAS controller and memory modules use even parity names while the SPC store uses odd parity names; consequently, diagnostic programs must be capable of distinguishing these two cases and proceeding accordingly.

2.2.4 Loading store from program tape unit

To save development effort, the SPC program used to load from the Program Tape Unit (PTU),[4] compare with PTU tape, and dump the memory contents of the SPC stores to a PTU tape is used to perform the same functions for the SSAS announcement stores. However, since there are some significant differences regarding tape format, loading procedures, and priority levels between the SPC stores and the announcement stores, some changes were made in the PTU program to extend its capabilities.

The requirements for the handling of the tape unit are as follows:

(*i*) The announcement information should be separate from office data and generic programs. A separate tape containing only announcement data should be used for announcement stores.
(*ii*) For the load operation, the craftsperson must force one SPC bus active, simplex, before using the program to access the announcement stores on the other bus.
(*iii*) The announcement tape header should be modified to distinguish it from an SPC tape.
(*iv*) SPC reliability should be given the highest priority. The announcement store should not be accessed if the SPC is not able to run full duplex with all stores operational.
(*v*) Changes must be made to minimize SPC downtime resulting from undetected SSAS store faults.

With these requirements in mind, the craftsperson is instructed not to use the program to access the announcement stores if the SPC cannot run full duplex with all stores operational. With the SPC fully operational, one bus may safely be forced active for the load operation, thus placing the system in the simplex mode. This choice of which bus to force active must be made such that the standby SPC bus will correspond to the SSAS side requiring PTU actions. This procedure ensures that the announcement stores will never be placed on the active SPC bus. Following this line of conservatism, the announcement stores are normally switched off the standby SPC bus and are connected just prior to a read or write instruction. After this instruction is completed, the stores are removed from the bus.

The SSAS side to be accessed must first be placed in the out-of-service or unavailable state for the compare or dump option. If the operation is to load an SSAS announcement store with data from tape, the SSAS side must be in the unavailable state. This state will not allow SSAS side switching if a fault should be detected on the active side.

2.3 SSAS fault recovery

2.3.1 PROCON detected and reported faults

SSAS fault recovery strategy is based on a fault detection scheme which uses both the SPC and the programmable controller (PROCON) to detect various faults associated with the SSAS hardware. Faults detected by PROCON are reported to SPC. The decision to choose one method or the other was based on the following considerations.

(*i*) SSAS faults, particularly those with serious impacts on call processing, should be detected as quickly as possible.
(*ii*) SPC routine processing time required for detecting SSAS faults should be kept as small as possible.
(*iii*) PROCON should not be responsible for detecting those faults which could affect its integrity or its ability to communicate with the SPC.

Those faults which are detected by PROCON are analyzed before being reported to SPC. PROCON has the capability to analyze available data and reinterrogate appropriate maintenance registers. If failures are indicated, the result of analysis and reinterrogation are sent to the SPC. The reports sent by PROCON are analyzed by the SPC and, if necessary, appropriate commands are sent to PROCON to prevent it from flooding the SPC with unnecessary information.

2.3.1.1 Continuous exercises. Continuous exercise of hardware by PROCON is a means of verifying the correct operation of the hardware and detecting various failures. Exercises are used both on active and standby SSAS sides, but these are particularly important on the standby side since no call processing is taking place. (Call processing exercises the SSAS rather completely.)

The announcement subsystem[5] (i.e., the announcement store and announcement distributor combined) is among those units which is continuously exercised by PROCON on both the active and standby sides to ensure the integrity of the speech heard by the customers. Any malfunction in these units would result in the appropriate flags being set in a maintenance register called the Error Summary Register (ESR). This register is periodically scanned by PROCON, which analyzes the content and sends the appropriate information to SPC.

2.3.1.2 Announcement subsystem failure reports. PROCON can recognize failures associated with the announcement subsystem by scanning the Error Summary Register (ESR) periodically. The presence of any announcement subsystem failure will result in one or more flags being set in the ESR. If the content of ESR is different from its last look, PROCON sends a message to SPC indicating what trouble has occurred, provided PROCON has received no commands from SPC to stop sending information.

During normal operation, a Hamming and parity check on the address and data received from the announcement stores is performed by hardware[5] before the data are sent on to the service circuits. Once an error is detected, an error flag will be set in the ESR. The 67 bits of address, data, Hamming and parity for the error will be trapped in special maintenance address and data registers.

The content of the ESR is scanned periodically by PROCON. Upon identifying an announcement store error flag in the ESR, PROCON saves the contents of trap registers. Then it performs a read using the saved announcement store failure address. Subsequently, it loads the FIFO output register with the available information. This information includes (see Fig. 10) the content of ESR, the content of failure trap address and data registers (both original and retry values), and an extended reply code used to distinguish the cases where no error is detected upon retry from those where error is also detected upon retry.

Upon receiving a PROCON report, SPC analyzes the reported errors and decides whether a hard fault or a transient error is present. Detected errors and faults are categorized according to their types and seriousness. Types of errors and faults detected by SPC are repeating double bit data fault, repeating single bit data fault, transient double bit data error, transient single bit data error, and repeating single bit address fault. The address and data associated with the errors and single bit data faults are also saved, and are printed hourly or upon manual request on the maintenance teletypewriter for manual troubleshooting. The failing address is also passed to the store diagnostic program to allow extensive tests of the suspected area of the memory.

A feature that will aid in detecting announcement store errors is the announcement store recital program, which sequentially accesses every announcement phrase by loading its address into the maintenance

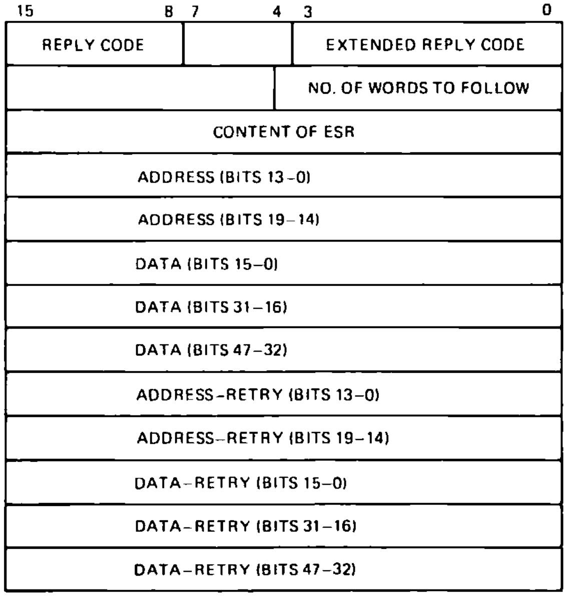

Fig. 10—Format of announcement subsystem failure report.

time slot of a Recirculating Shift Register (RSR). This type of routine exercise will guarantee that the memory used by every speech segment is accessed periodically. If any store error associated with a particular speech segment exists, a flag will be set in ESR notifying PROCON of the error and the address and data associated with the errors will be trapped. This program is implemented on both active and standby SSAS sides.

2.3.1.3 Mate frame buffer failure reports. Communication between the two SSAS sides (see Fig. 11) takes place asynchronously through a first-in, first-out serial memory called the Mate Frame Buffer (MFB). Each transmission through the MFB usually consists of data and an ID field which identifies the data.

During normal operation, all the data associated with the calls in progress are sent from the active side to the standby side through the MFB. This operation enables the standby side to have an up-to-date copy in its Random Access Memory[5] (RAM) of the data associated with each call for use if it is required to change role from standby to active. To test the validity of the data before it is stored, a 1-out-of-16 encoded ID word is associated with the data coming through MFB, to indicate what it contains. The standby PROCON checks for a valid ID format and for proper sequences of IDs. If the standby PROCON detects an error, it informs the SPC by loading appropriate data in its FIFO output register.

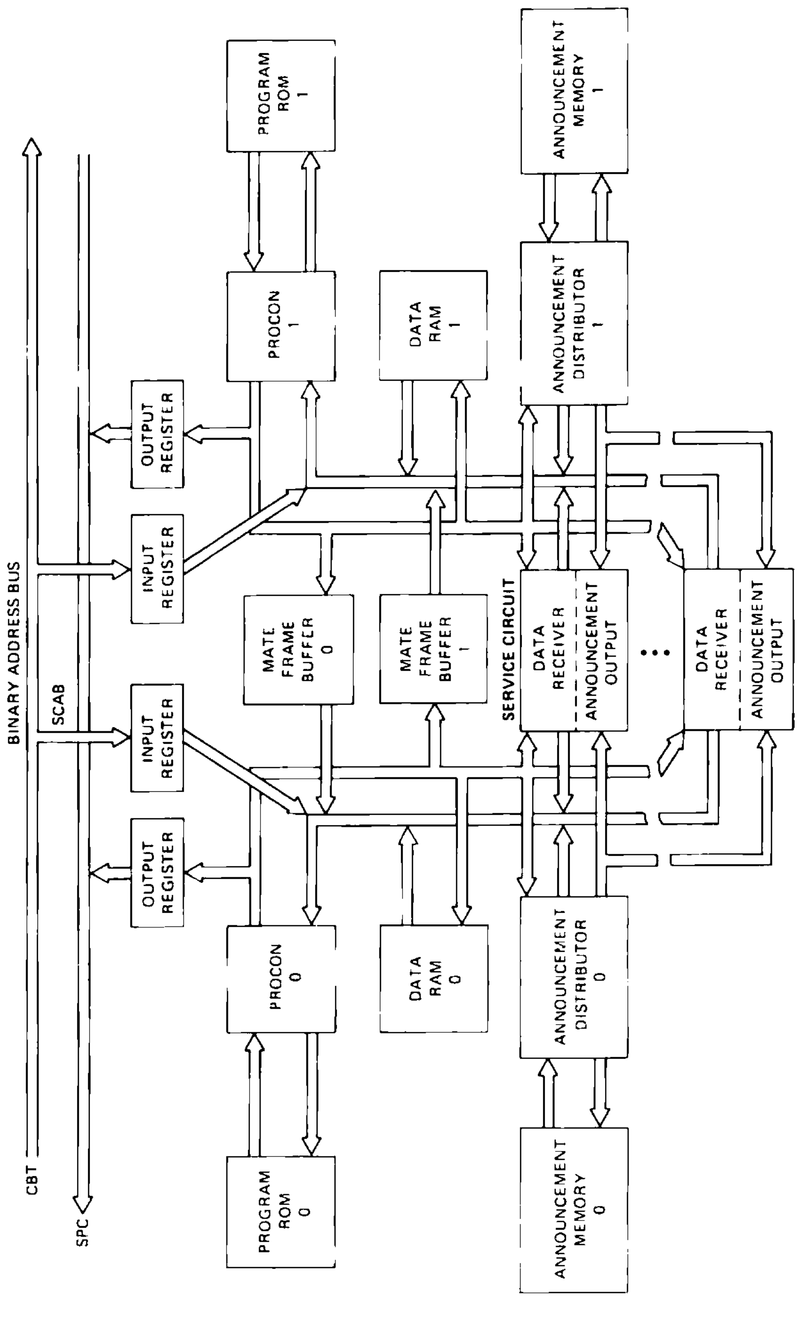

Fig. 11—Duplicated SSAS configuration.

During the unloading of the MFB by the standby side, an overflow condition in the MFB results in an MFB full flag being set in the ESR associated with the active side. This overflow condition could be caused by either an MFB failure or a failure in the standby side which prevents it from unloading the MFB. The MFB full flag in the ESR is monitored by the active PROCON. Upon detecting the MFB full condition, the active PROCON notifies SPC by loading the necessary information in the FIFO output register.

2.3.1.4 CDA error reports. A CDA failure is recognized by PROCON during normal interrogation. All calls except those that are in the announcement mode only require interrogating the CDA reply register for coin deposits. Prior to sending the first interrogation command to a CDA, PROCON reads the reply register to ensure that there are no premature replies from the CDA. In the case of premature reply, all records of the call are erased and a message is sent to SPC indicating a bad CDA. During the course of interrogation, the parity of the data received from the CDA is checked by PROCON. In the case of a parity failure, PROCON again erases the records of the call and informs the SPC of the failure.

During the course of interrogation, PROCON continuously monitors the format and the validity of the data in the CDA reply register. A reply with invalid format is considered a CDA failure. In this case, the records of the call are also erased, and SPC is notified of the CDA failure. Another means of CDA error detection is the CDA maintenance register. If a maintenance bit is set in the reply register, the maintenance register is read to check the type of error. In this case again, PROCON notifies SPC and erases the records of the call. Upon receiving a CDA error report from PROCON, SPC routes the call to an operator position and also takes the bad CDA out of service, so it is not used for any future call. A diagnostic automatically will be requested for the bad CDA to aid the craftsperson in locating the faulty units. If too many CDA failure reports are received within a certain period of time, all associated with the same group controller, SPC then assumes that the group controller associated with the active side is faulty.

PROCON can only detect failures in the digital circuitry of a CDA. Failures in the analog circuitry of a CDA are detected by an error analysis program which resides in SPC. The success rate of attempted calls for a given CDA is analyzed, and if it is below a certain average threshold, the CDA is considered faulty and is taken out of service and automatically diagnosed.

2.3.2 SPC detected and reported faults

2.3.2.1 Maintenance interrupts. Maintenance interrupts are used as a means of identifying faults in SPC-SSAS interface registers. These types of faults have a serious impact on the operation of SSAS. They can result in isolation of the SSAS, thus crippling the capability of SPC to determine the basic integrity of the SSAS. This basic integrity is a necessary condition for the SPC to have confidence in the capabilities associated with PROCON in detecting SSAS faults.

As with other TSPS peripherals, any malfunction in SPC-SSAS communication is detected by using standard peripheral unit fault recognition techniques that can cause an F-level interrupt (such as central pulse distributor enable verify, and all-seems-well signal failures).[4] Upon occurrence of an F-level interrupt, appropriate fault recognition programs are entered to analyze the faults present in SPC-SSAS interface areas. These programs have the capability of distinguishing serious faults from nonserious faults and reporting them to appropriate reconfiguration programs.

During a scan operation of SSAS by SPC, a malfunction in one or both Scan Answer Buses (SCAB) results in an SPC processor mismatch which in turn causes a C-level interrupt to occur.[4] Upon occurrence of a C-level interrupt, appropriate fault recognition programs are entered which first determine whether SSAS caused the interrupt, and then by performing some tests isolate the faulty SCAB. The nature of failure and its degree of seriousness are also passed to SSAS reconfiguration programs.

2.3.2.2 Periodic scans. SPC detects faults in PROCON itself or in the PROCON peripheral control system by interrogating maintenance registers on the scan answer bus. This interrogation is done periodically (every 25 ms). The information provided by this interrogation (scan) is used by the SPC to determine the sanity of the PROCON for processing calls and to verify the capability of PROCON to detect SSAS faults. Among faults detected by periodic scanning are PROCON all-seems-well failures, Random Access Memory (RAM) all-seems-well failures, PROCON clock-stopped and output register-full indications. Both active and standby SSAS sides are monitored by the periodic scanning. The scan result is analyzed and is used to distinguish serious faults from nonserious faults, and to determine whether or not an immediate action is needed by the SPC.

2.3.2.3 Dead side tests. Despite the fact that extensive error checking circuitry has been built into the SSAS hardware, the possibility still exists that some faults may elude the checks and not be reported to the SPC. Potential faults in this category are nonclassical faults, multiple faults, certain PROCON and RAM failures, etc. A possible consequence of such faults is a dead SSAS side which appears to be free of faults during a period of no normal SPC-SSAS communications. To detect such a failure mode, the SPC performs an exercise periodically. The exercise involves sending a command to an SSAS side and receiving a specific reply. If an SSAS side fails to reply as expected, then the side is considered to be faulty, and appropriate information is reported to the reconfiguration program.

Since SSAS functional integrity is constantly checked by call processing functions, the dead side test for an active SSAS side will be skipped if traffic is present on the side. More specifically, an SSAS side will not be tested if valid replies are coming through its output register at a regular interval.

2.3.3 Reconfiguration

2.3.3.1 Classification of faults. The response of the SSAS reconfiguration program to various SSAS faults is based on a fault classification scheme which categorizes the faults according to their impact on the normal operation of SSAS. Faults fall into one of four categories which can be summarized as follows:

(*i*) Serious faults. This type of fault corresponds to those which seriously affect the call processing capability of SSAS. For example, a fault which stops normal operation of PROCON is considered to be serious. A serious fault on an active SSAS side will cause an immediate switch to the other side.

(*ii*) Nonserious faults. A nonserious fault is defined as a type which does not seriously affect the call-processing capability of an SSAS side. An active side with this type of fault can continue to perform its normal function adequately as long as it is required to do so. This in turn implies that the switching action to the mate side can be delayed indefinitely, if necessary. For example, a single bit failure on the announcement data can be considered nonserious since its impact on the quality of the announcement is insignificant. As another example, a group controller failure on the active side can also be considered nonserious, since the side can continue to process calls using the remainder of the group controllers and their associated CDAS.

(*iii*) MFB faults. This type of fault is usually an indication of a failure in the MFB itself or the standby SSAS side, with no impact on the active SSAS side. No switching of sides is performed in this case. Instead, the standby side and MFB are both diagnosed for isolation of faulty units.

(*iv*) CDA faults. This type of failure usually represents a problem in an individual CDA. Since there are many CDAS in SSAS, an individual CDA failure does not have a serious impact on the call-processing capability of SSAS. A faulty CDA is marked out of service and will not be used until it is diagnosed and repaired.

2.3.3.2 Smooth side switch. The process of interchanging the active side and standby side is referred to as a side switch. If this is done so that the flow of announcements is not interrupted, the switch is called a smooth switch.

A smooth switch between two SSAS sides is performed under two circumstances. One case is when the active side has a nonserious fault such as single bit failure in announcement data. In this case, the system can afford to wait as long as needed to switch smoothly. The other situation is when a routine switch to the standby side is desired to test the hardware controlling the reconfiguration and to provide a chance to run a routine set of diagnostics on the previously active side. This process (the smooth switch exercise) is performed three times a day.

A smooth switch is typically characterized by the following actions. First, the SPC sends a command to the active side, telling it to go standby gracefully. Second, the active side brings the standby side in loose synch with the active. Third, the standby side is told to go active gracefully. Since by now the active and the standby are loosely in synch, both sides time themselves until the right moment when the old active goes standby followed almost immediately by the old standby going active. The new active verifies its success by sending a reply to the SPC. Upon receipt of this reply, the SPC changes the SSAS status words to record the new configuration. This type of side switch does not provide any discontinuity in the announcement heard by the customer and also has no effect on the call processing operation.

2.3.3.3 Other types of side switch. The handshaking between the two sides required to perform a smooth switch may not be possible if the active side develops a serious fault such as a PROCON all-seems-well failure. Assuming that the nonactive side is in the standby ready state, the SPC simply takes the faulty side out of service and makes the standby side active. This type of switch is referred to as an *immediate switch*. Since the standby's data RAM is up to date with the active side, only a minor disruption to call processing results. The announcement heard by a customer is interrupted. Approximately a half-second elapses, and then the entire announcement is repeated. In addition, it may result in mishandling a single call. However, this minor disruption of the call processing is considered a small penalty, as hardware faults occur rarely.

Another type of side switch, called a *rough switch*, happens when the active side is doing call processing and the nonactive side is running a routine diagnostic on itself. If the active side develops a serious fault, it is immediately taken out of service. In the meantime, the routine diagnostic on the nonactive side is aborted, and the side is made active. Since the data RAM on the standby side is not up-to-date while a diagnostic is being run, it will be initialized before becoming active. Therefore, all processing on existing calls is aborted and the SSAS starts afresh. Considering the reliability objectives of the SSAS, the probability of occurrence of a rough switch is estimated to be very small.

2.3.3.4 Tolerance of software errors. Software errors are usually analogous to a serious component failure, causing both SSAS sides to go down and remain down until the problem has been rectified. However, certain classes of software errors could be considered tolerable. An erroneous software instruction which is executed infrequently, if it is considered to be nontolerable, could cause a permanent outage of the SSAS upon its first execution. On the other hand, if at this time the SSAS is reinitialized and restarted, it will continue to perform its normal function until some later time, when it will encounter its second failure, due to the execution of the erroneous software instruction. Following this strategy would substantially decrease the SSAS downtime in cases where infrequently executed software errors are present.

The above considerations provided the basis for developing a strategy which tolerates certain classes of errors in the PROCON software. Upon occurrence of an SSAS outage, the SPC performs an analysis for both sides using past history of SSAS outages. First the side with the fewer past failures is selected. Then after the analysis a decision is made on the nature of the problem for that side, as to whether a hardware fault or a software error exists. In the latter case, the side is reinitialized and used for call processing after a successful initialization, while a diagnostic is requested for its mate side. In the former case, both SSAS sides are diagnosed to pinpoint the hardware malfunctions.

2.3.3.5 Duplex failure. If SSAS becomes totally unavailable, the call processing function performed by the SSAS is aborted and all new calls will be routed to the operators. The new calls are blocked from attempting to cease CDAS for service. Also, several clean-up actions are performed by the SPC to allow for rapid restoration of ACTS. For example the peripheral orders waiting to be sent to the SSAS are searched and disposed of. The SSAS is blocked from providing ACTS until all the clean-up tasks are completed.

2.4 SSAS diagnostics

2.4.1 Controller diagnostic

2.4.1.1 Interface SPC—PROCON. The SSAS diagnostic programs run on one SSAS side at a time, always a nonactive side. The diagnostic request is made either by manual request or by fault recognition programs, either as a result of a fault being detected or as a routine diagnostic. The routine diagnostic is run once a day to do a complete check of the SSAS hardware.

As a direct consequence of the SSAS architecture, the SSAS diagnostic control is divided into two parts; the first part resides in SPC, and the second resides in PROCON's Read Only Memory (ROM) and is under the control of a standby monitor program which controls the various activities of a nonactive SSAS side, including diagnostics. Other activities associated with the standby monitor include fault recognition, unloading of the Mate Frame Buffer (MFB) and other SPC-requested nondiagnostic functions. SPC retains the overall control over the SSAS diagnostic. SPC initiates SSAS diagnostics by first marking the SSAS side out of service and then sending a diagnostic request to the PROCON. The request indicates that a group of tests referred to as a "phase" is to be run by PROCON. In some cases, this request is also accompanied by a test pattern to be used by PROCON. The standby monitor decodes the request, performs error checking, and invokes the PROCON diagnostic control program. This program further decodes the diagnostic request to select the proper diagnostic phase. Once a diagnostic phase is initiated, the standby monitor exercises control over the execution of the diagnostic. A diagnostic phase performs tests on the SSAS functional block being diagnosed, using predetermined test patterns and sends the test results to SPC through the FIFO output register. The test results are used by the SPC to isolate the particular fault involved.

2.4.1.2 Levels of raw data compression. The diagnostic test results sent by PROCON through the FIFO output register are compared against the expected result by the SPC. The result of that comparison is saved as a binary string of ones and zeros referred to as "raw data" results. The zeros represent the tests passed and ones indicate the tests failed. Although PROCON does considerable data comparison, in many diagnostic phases the test results sent by PROCON are quite voluminous and exceed the capacity of the trouble number generation programs. In such cases, SPC uses a compression algorithm for converting the test results into a compressed raw data bit string. The algorithm uses a simple procedure: If any test result word sent by PROCON does not match the expected value, the SPC diagnostic program will map this mismatch into one raw data bit. The compressed raw data bit string is further reduced by the SPC to a 4-part, 16-digit number referred to as a "fault signature," which uniquely identifies the string. The fault signature is used by the maintenance personnel to reference a trouble-locations manual to isolate the particular fault involved.

In some cases, the fault signature generated by SSAS diagnostic programs fails to provide enough information to clear the trouble associated with the SSAS. In such cases, the maintenance personnel can manually request either compressed raw data or expanded raw data. The raw data printout is used by maintenance personnel as an aid in manual troubleshooting.

2.4.1.3 Power of PROCON test. The SSAS architecture has partitioned the hardware in functional modules (such as PROCON, peripheral control, announcement memory, service circuit group controllers, etc.).[5] Each of these functional modules are further partitioned to attain the desired level of maintainability and diagnostic resolution.

PROCON access to each functional module results in a simple diagnostic structure consisting of segments that diagnose only one functional module. In general, diagnostic resolution is directly related to the size of the module's diagnostic. Resolution can be improved by reducing the size of the module. A further improvement can be realized by choosing a diagnostic structure that begins with the diagnosis of the most basic functional module and proceeds to test each module by adding diagnostic segments that use the diagnosed modules. These considerations have been taken into account in planning the SSAS diagnostic structure, which provides an average diagnostic resolution of three or fewer circuit packs, with the capability of detecting greater than 90 percent of potential SSAS faults. These diagnostic objectives are required to limit the average repair time to about 2 hours as specified by the SSAS maintainability objectives. PROCON access capability combined with the modular architecture allows the SSAS diagnostic objectives to be achieved.

Another feature that helps to achieve the SSAS diagnostic resolution objective is the bit-sliced architecture used in designing SSAS input and output bus interfaces. The bit-sliced architecture is used to aid the SSAS diagnostics in pinpointing faulty packs and to reduce the number of board codes. This means that the common bits for all registers on a bus are grouped together on one or a few circuit packs (i.e., bit *i* for all registers on the bus are grouped together). Because of this architecture, it is only necessary for the SSAS diagnostic to know which bit is bad to isolate the faulty circuit boards.

2.4.2 CDA diagnostic

2.4.2.1 SPC-PROCON interaction. The CDA diagnostic is initiated by the SPC either as a result of CDA failures detected by fault recognition programs, error analysis programs, or a manual diagnostic request via the maintenance teletypewriter channel. Also, an automatic progression test of the CDAS is performed to ensure that each CDA is diagnosed once a day.

A service circuit known as the CDA test circuit is used to diagnose CDAS. It has network appearances so it can be connected to the various CDA ports for analog circuit tests.[5] When a CDA diagnostic is requested, the SPC attempts to establish a test configuration if the requested CDA is available for testing. It first verifies that the CDA test circuit is available and the nonactive SSAS side is in the standby ready state. If not, the diagnostic is blocked; otherwise, the CDA test circuit is made maintenance busy, and orders are sent to the active SSAS side to switch both the CDA test circuit and the CDA under test over to the standby SSAS side. Next, the Signal Distributor (SD) controlled relays in both the CDA and CDA test circuits are released without regard to their previous states, and the two circuits are connected via the TSPS network.

Once the test configuration is established, various phases of the CDA diagnostic are initiated by the SPC. For each phase, the SPC first operates the required SD-controlled relays, and then sends an order to the standby SSAS side to run tests of the specified phase.

The PROCON standby monitor program processes the diagnostic order and transfers PROCON control to the CDA diagnostic program. This program performs the tests of the specified phase and sends the results to the SPC via the FIFO output register. Once the PROCON reply is received, the SPC compares the test results with the expected results. The result of this comparison is the diagnostic raw data. Each bit set to a one in the raw data represents a test failure. The raw data is then converted into a trouble number which is used by the maintenance personnel to isolate the trouble.

When all the phases of the CDA diagnostic have been completed, the SPC releases the relays of the CDA and CDA test circuits, breaks the network connection between them, and sends orders to the standby SSAS side to switch them back to the active side. The CDA test circuit is returned to the maintenance idle state, and the CDA under test is either returned to service or placed out of service, depending upon whether the diagnostic tests have been passed or failed.

2.4.2.2 CDA test circuit use. The CDA test circuit is a special type of SSAS-controlled service circuit. There is one CDA test circuit per SSAS, and it is always equipped as circuit 15 of the service circuit frame 1.

The primary functions of the CDA test circuit can be summarized as follows:

(*i*) It produces simulated coin deposit signals for testing a CDA coin tone receiver.
(*ii*) It detects test tones produced by a CDA announcement decoder.

The CDA test circuit architecture is quite similar to that of the other SSAS service circuits. It consists basically of digital circuitry to interface with SSAS, a tone generator and a tone detector (see Fig. 12). Its digital interface circuitry with SSAS consists of command decoding logic, reply registers, and control logic which provides additional decoding of CDA test circuit commands, and also controls the miniature relays which provide switching of pads and filters.

The tone generator consists of an announcement decoder, an inter-

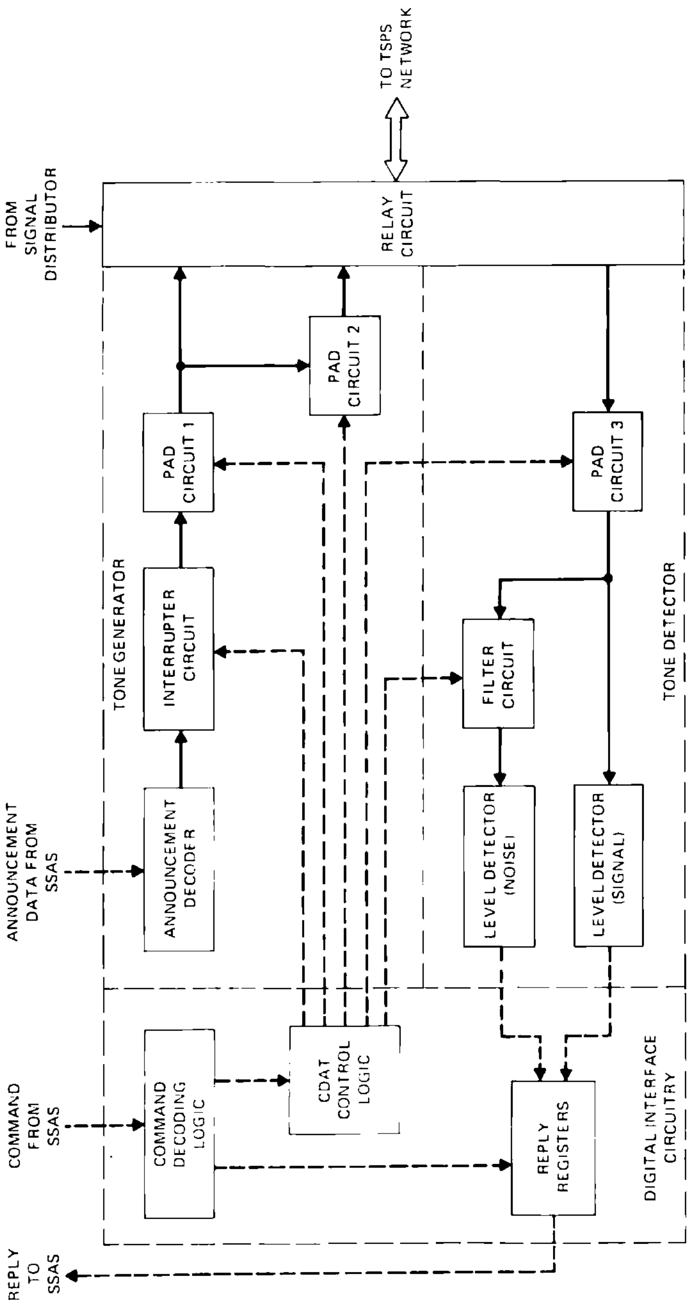

Fig. 12—Block diagram of CDA test circuit.

rupter switch, and a set of switchable attenuation pads. The interrupter switch is used to clamp the analog output of the announcement decoder to ground, thereby silencing the tone generator. The switchable pads provide the level variation needed for a CDA coin tone receiver testing.

The tone detector consists of a switchable pad circuit, level detectors, and a switchable filter circuit. The switchable pad circuit provides level adjustment for signals generated by a CDA announcement decoder. The output of the announcement decoder of a CDA circuit is tested by the level detectors for both signal loss and distortion. During the process of checking for signal distortion, the switchable filter circuit is used to remove the expected fundamental frequency component of the announcement decoder test tone without attenuating any harmonic frequency components.

2.4.3 Use of test tones from announcement store

Many tests performed during the course of CDA diagnostic are aimed at detecting possible malfunctions in either the announcement decoder or the coin tone receiver associated with the CDA under test. Proper operation of the announcement decoder is verified by monitoring its response to a set of digitized test tones. This response is checked for possible abnormalities in both frequency response and output level using a tone detector associated with the CDA test circuit. The coin receiver is tested by using the tone announcement decoder of the CDA test circuit to produce coin deposit signals from data stored in the SSAS announcement memory, and checking the response of the coin tone receiver to deposit signals having various level, frequency, and timing characteristics.

The test tones required for checking the proper operation of the announcement decoder and the coin tone receiver are stored as announcement data. A total of 18 test tones are required for CDA diagnostic tests. The storage requirement for many of these tones is significantly reduced by packing several of them into a single announcement phrase. Combining up to four 125-ms test tones into a single segment reduces the number of the required segments to six. One hundred twenty-five milliseconds allows ample time for the announcement decoder output to stabilize and for the CDA test circuit tone detector to achieve a stable measurement. However, this approach adds additional timing functions for the standby PROCON to perform. PROCON must determine the appropriate time to strobe the tone detector and, for some tests, it must blank out the unwanted tones by operating an interrupter switch within the CDA test circuit which turns off the analog output of the tone announcement decoder.

2.5 CDA alignment and manual tests

2.5.1 Trunk test panel connections

The TSPS Control, Display, and Test (CDT) panel provides a means of manually initiating tests on trunks, operator positions, and service circuits appearing on the TSPS switching network.[6] In conjunction with the ACTS feature, provisions have been made for testing both the CDA and the CDA test circuits from the CDT panel.

Most CDA circuit tests originating from the CDT panel require that test connections be established between the CDA circuit and the CDT circuit. Operation of the Master Test Line–Test key on the CDT panel causes the SPC to connect an MF receiver to the Position Master Test Line (PMTL) of the CDT circuit via the TSPS network. Once this connection is established, the SPC lights the Master Test Line (MTL) lamp on the CDT panel, indicating that codes can be entered via the CDT multifrequency (MF) key set. At this point, a test connection can be requested with any CDA circuit appearing on the TSPS network by entering its trunk group and member number and the camp-on bit via the CDT MF key set. A special display on the CDT panel indicates to the maintenance personnel whether the request was successful, aborted due to system failure, or invalid.

The SPC indicates the status of the requested CDA on a special display lamp. This lamp indicates whether the CDA is idle, traffic busy, maintenance busy, or out of service. If the CDA is busy and camp-on is requested, the SPC reserves it for maintenance by placing it out of service as soon as it becomes idle and notifies the maintenance personnel of its availability via a teletypewriter output message. If the CDA is idle or out of service, the SPC disconnects the MF receiver from the PMTL, extinguishes the MTL lamp, and establishes the test connection between the CDA and CDT circuits.

Several other test features involving CDT panel key actions have also been provided for CDA circuits. By operating appropriate keys, the maintenance personnel can place the CDA connected to the CDT out of service, or return it to service. Also, a display of the ferrod states of the CDA under test can be requested by operating a special key. This key action causes the scan row containing the CDA ferrods to be displayed on the program display of the control and display panel.

2.5.2 Trunk test panel measurements

The CDA circuits contain analog transmission components which provide a part of the talking path between the customer or customers and the CDA announcement decoder. These CDA circuit voice paths must be tested with respect to noise and signal loss whenever transmission problems associated with the circuit are suspected. The trunk test panel, shown in Fig. 13, provides a convenient means of performing noise and loss tests on CDA circuits that appear on the TSPS network.

Testing of the transmission path of a CDA circuit involves separate tests for performing noise and loss measurements. First, a test connec-

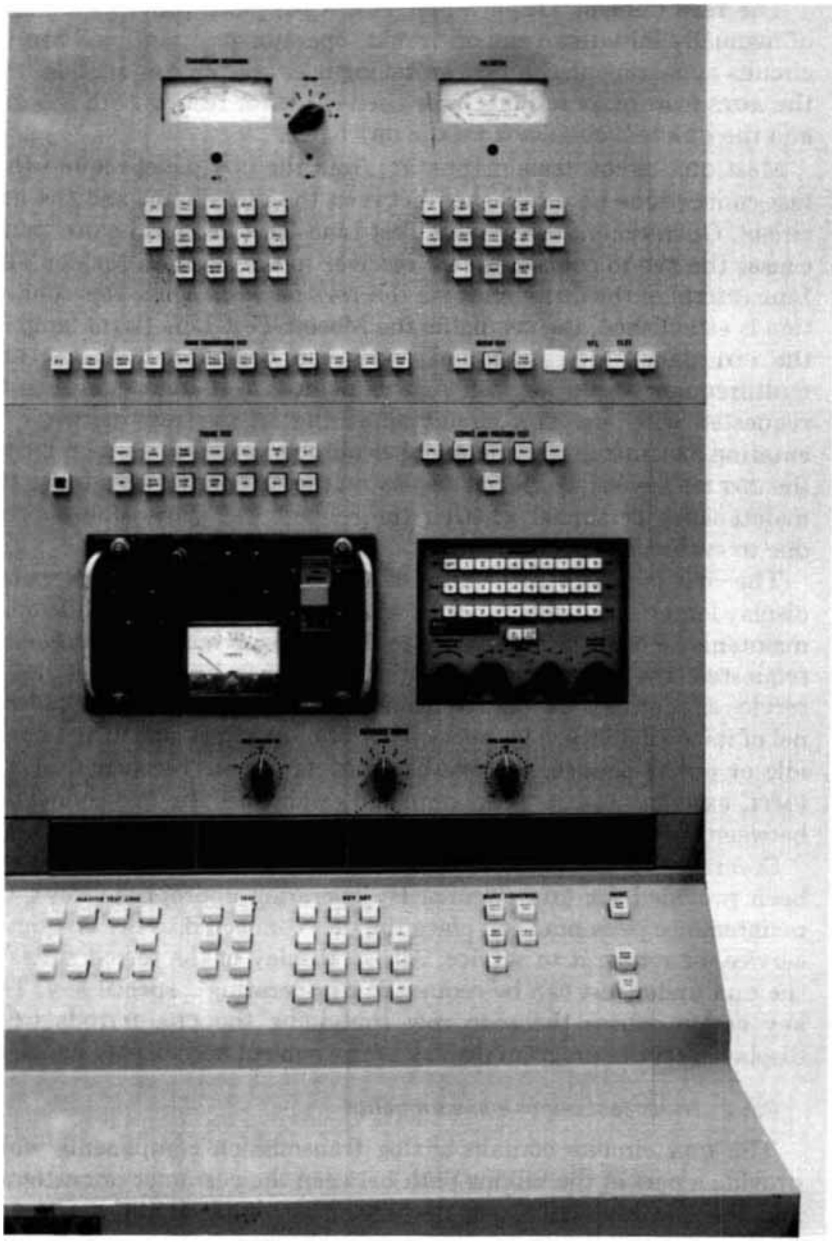

Fig. 13—Trunk test panel.

tion is established and the Master Test Line–Test key is operated to signal the SPC to connect an MF receiver to the PMTL. Next, by entering an appropriate test code and appropriate test control key actions, each of these tests can be initiated. These tests can be repeated or performed in any sequence as long as only one test control key is operated at a time. These tests can also be sequenced under the SPC control by using an appropriate test code. As soon as the code is entered, the SPC establishes the connection for the first test. Upon completion of the first test, the connection for the second test is established. This process is continued until the completion of the last test. At this point, the CDA circuit is disconnected from the CDT panel and is returned to its previous state.

Manual transmission tests of all CDA circuits in the office can be initiated by entering a special test code. The SPC connects each CDA circuit in sequence to the CDT panel. If the circuit is traffic or maintenance busy, it is skipped and a teletypewriter message is printed. After the completion of the tests, the SPC disconnects the CDA circuit from the CDT and connects the next one.

Automatic sequencing of the transmission tests of all CDA circuits can be initiated by entering another special code. The SPC connects each available CDA circuit to the CDT panel, and sequences through each of the tests. Once all of the tests for a CDA circuit are complete, another CDA circuit is connected. Busy circuits which are skipped are identified by teletypewriter messages.

2.5.3 Vocabulary recital

Because the announcement decoding aspect of the SSAS has no existing counterpart in the TSPS, a totally new test feature is provided from the CDT panel. This feature is an announcement vocabulary recital exercise whereby the maintenance personnel can listen to each speech segment in the announcement vocabulary via the CDT telephone head set.

Listening to the announcement output of a CDA circuit provides a manual check not only of the CDA circuit announcement decoder, but also of the SSAS announcement subsystem circuits. When an announcement store is loaded from tape, this feature can be used to verify the contents of the store before it is made active. The announcement recital test can be performed from both active and standby SSAS sides. Two options are provided for the recital test. They can be summarized as follows:

(*i*) The decoding of every announcement segment and phrase in sequence.

(*ii*) The repetitive decoding of a specified announcement segment or phrase.

The announcement recital test on an active SSAS side begins with a CDA circuit connected to the CDT panel. By operating appropriate keys and entering a special test code, the SPC starts the PROCON-controlled announcement exercise on the active side. This exercise accesses the announcement data for a phrase and sends it to the CDA circuit under test. It sequences through every word or phrase in the announcement vocabulary and then stops. Each word or phrase is separated by ½ second of silence.

Entering another test code causes a specified announcement phrase to be repeated at a constant rate. This is controlled by another PROCON exercise which accesses the announcement data for the specified phrase and sends it to the CDA circuit under the test. One-half second of silence is inserted between repetitions. The phrase is repeated until the circuit is released from the CDT panel.

Similar procedures are used for the announcement recital test from a standby SSAS side. Again by entering special codes, appropriate announcement exercises are started. If the SSAS diagnostic is in progress on the standby side, the exercise is not run and a teletypewriter output message is printed.

III. ACKNOWLEDGMENTS

The authors wish to acknowledge the contributions of many colleagues who took part in the development of both the operational and the maintenance software described in this article. Among them, several people deserve special mention: J. Atkins, J. R. Connet, D. C. Dowden, I. S. Dowden, R. V. Miller, and R. D. Nafziger for their work in operational programs; M. W. Medin, D. G. Raj-Karne, W. R. Serence, and M. D. Soneriu in maintenance software; and R. H. McGuigan and S. B. Windes in growth and retrofit.

REFERENCES

1. M. Berger, J. C. Dalby, E. M. Prell, and V. L. Ransom, "TSPS No. 1: ACTS Overall Description and Operational Characteristics," B.S.T.J., this issue, pp. 1207-1223.
2. A. W. Kettley, E. J. Pasternak, and M. F. Sikorsky, "TSPS No. 1: Operational Programs," B.S.T.J., *49,* No. 10 (December 1970), pp. 2625-2683.
3. 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.
4. G. R. Durney, H. W. Kettler, E. M. Prell, G. Riddell, and W. B. Rohn, "TSPS No. 1: Stored Program Control No. 1A," B.S.T.J., *49,* No. 10 (December 1970), pp. 2445-2507.
5. G. T. Clark, K. E. Streisand, and D. H. Larson, "TSPS No. 1: Station Signaling and Announcement Subsystem; Hardware for Automated Coin Toll Service," B.S.T.J., this issue, pp. 1225-1249.
6. W. K. Comella, C. M. Day, Jr., and J. A. Hackett, "TSPS No. 1: Peripheral Circuits," B.S.T.J., *49,* No. 10 (December 1970), pp. 2561-2623.