THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 61, No. 5, May–June 1982
Printed in U.S.A.

1A Voice Storage System:

Office Engineering, Maintenance, and Reliability

By P. W. BOWMAN, J. C. KENNEDY, and G. A. VAN DINE

(Manuscript received July 31, 1979)

The quality of service that will be delivered by the Voice Storage System (1A VSS) is influenced by a number of diverse factors which are addressed in this paper. High intrinsic reliability is designed into the system at every level; it is built from system redundancies, defensive software strategies, hardware self-checking, and manufacturing quality control. Good service also requires a complement of equipment which can adequately handle peak traffic loads; the sizing of this complement is the job of Office Engineering. Effective maintenance software will minimize, and in most cases obviate, the impact on service of circuit component failures. And finally, effective plans and test equipment are required to ascertain that a newly installed 1A VSS properly performs its intended functions.

I. INTRODUCTION

The 1A Voice Storage System (1A VSS) comprises a new functional node in the telephone network, which will provide a number of new services classified generally as Custom Calling Services II (CCS II). These services provide the means for conveying voice messages between customers who do not happen to be at their telephones at the same time.

A conventional telephone switching office mediates communication between two parties by setting up a connection over which they may converse. Unlike a switching office, the 1A VSS takes on the role of one of the conversing parties. In the course of handling a typical call, the 1A VSS receives input control signals at various times, outputs system announcements in response, and either records or retrieves customer messages.

The quality of service is important; yet, this type of communication is complex; there are many opportunities for system malperformance to disrupt the call. The following paragraphs introduce several factors that contribute to high quality service; they are expanded upon in the remainder of this article.

Office Engineering, which is covered in Section II, is the process through which a Bell operating company (BOC) decides specifically what equipment to order to provide CCS II.[1] Inherent in the process are estimates of the amount of traffic to be expected as a function of time. This information is derived from market studies which project the expected usage of specific services in specific customer sectors.

Once a system is put into service, however, it is important to obtain feedback which will allow the early usage estimates to be examined in light of actual market acceptance. With such information, intelligent decisions can be made for office growth, which tracks increasing service demands. The collection and utilization of these usage data is covered under Traffic Measurements in Section 2.4.

With 1A VSS services in use, the system is being entrusted with large numbers of customers' messages. Extensive precautions have been taken in all phases of the system design and manufacturing process to keep these messages secure and to assure their delivery. The design and manufacturing strategies employed to assure dependable 1A VSS performance are covered in Section III under the heading of Reliability.

Although the best quality obtainable with present technology is built into the 1A VSS components, precautions are taken against the occurrence of component failures by providing spare equipment and maintenance software which can locate the malfunctioning component. This software then responds to system troubles by configuring a working system that excludes the faulty unit. The philosophies and methodology employed to keep 1A VSS operational and to facilitate repair are presented in Section IV.

As a final step in office installations, many system tests are performed by the Western Electric installation organization. These are followed by BOC tests which verify the system's interfaces with other systems. Associated feature tests verify that the 1A VSS works properly when accessed through the connecting ESS offices. This testing is covered in Section V.

II. OFFICE ENGINEERING

2.1 Planning

The engineering of a 1A VSS office and of the associated equipment in 1/1A ESS is performed to determine the specific equipment that is required to provide CCS II. Initially, trial office engineering is performed for economic feasibility study purposes. Such trial economic studies result in a decision as to whether or not to proceed with a local offering of CCS II. After a decision is made to proceed with 1A VSS, an order is placed for the required equipment.

To engineer a 1A VSS office and the equipment in the associated 1/1A ESS offices, CCS II marketing information is necessary. Estimates of the rate of penetration and the ultimate market penetration for the services being provided are required for three types of stations: residential, small business, and large business. Providing this essential input is the responsibility of the BOC marketing organization.

The schedule for introducing CCS II in the 1/1A ESS offices must also be determined. For each 1/1A ESS that will provide these services, the number of residential, small business, and large business stations is required. The estimated annual growth rate for each of these categories, and the period of time for which the equipment is being engineered, must also be specified.

The above marketing and 1/1A ESS information comprise the required inputs to the worksheets which guide BOC engineering personnel through the numerous calculations required to determine the amount of equipment required to provide CCS II services and to support their growth.

2.2 Storage and offered load

Storage time is a new traffic parameter associated with 1A VSS. Storage time is defined as the time a voice greeting or message is stored in 1A VSS before it is erased. The average storage time multiplied by the number of greetings and messages recorded during this time gives the number of simultaneous greetings and messages in the system. This, in turn, multiplied by the average length of a voice greeting or message determines how much storage is required, and hence the number of disk transports in a 1A VSS office.

The two main parameters that must be determined to engineer a 1A VSS office are (*i*) the offered load (erlangs) between each 1/1A ESS and 1A VSS, and (*ii*) the total storage (disk capacity) required in the 1A VSS. These are calculated from numerical arrays which contain market penetration estimates and traffic characteristics tabulated by service and customer types. The two-way trunks between 1/1A ESS and 1A VSS are engineered for 0.01 probability of blocking. The disk transports are engineered for more than an order of magnitude better blocking performance than the trunks.

2.3 Equipment engineering

For the No. 1/1A ESS offices, CCS II services require new dial pulse repeating/monitor trunks, additional program store and additional call store. The increase in program store is required to hold the CCS II feature; the additional call store holds the service translations. Since additional ESS capacity is required for 1A VSS-handled calls and for advance calling delivery, some 1/1A ESS processor capacity must be allocated for these services. The CCS II services have a minimal effect on the 1/1A ESS network and service circuits.

The 1A VSS office equipment[2] (see Figs. 1 and 2) can be divided into three categories. The first category is comprised of a minimum basic set of common equipment required for all offices. Included in this first category are the two Auxiliary 3A Processors (labelled generically as the Central Processor in Fig. 1), main memory storage, the two peripheral controllers, three storage media controllers (SMCS) and a service circuit frame.

The second category covers major frames or units of equipment which are either traffic sensitive themselves or which must be provided to support traffic-sensitive equipment. Frames or units in this category must be provided by the supplier during the system installation, or at major growth periods. Typically, this equipment is engineered for a two-year period. Included in this second category are voice access circuit frames and units, SMCS, and uninterruptible power equipment called triports. The voice access circuit frames in this category commonly contain additional unequipped circuit pack locations to accommodate equipment in the third category.

The third category is comprised of traffic-sensitive plug-in equipment, which includes trunk access and buffer circuit modules, plus disk transports. These units are procured as required and may be installed and turned up for service by the BOC maintenance personnel. Such growth results from adding 1/1A ESS connecting offices, penetration into the potential user market anticipated in the market forecast, or a rise in actual trunk group and storage usage based upon traffic measurements. Equipment in the third category is typically engineered for a 6-month period.

2.4 Traffic measurements

Traffic measurements are provided on periodic reports referred to as C, H, Q, and D schedules. The C schedule lists measurements of trunk utilization, while the H schedule provides measurements of internal subsystem utilization. Measurements which are taken to reflect the load on internal subsystem resources include counts of disk transport, service circuit and disk subsystem usage, total calls handled, processor real time, processor main memory usage, and a record of overload control actions. Both the C and H schedules are half-hour reports. The quarter-hourly Q schedule is a subset of the H schedule that provides a quick look at a few key measures of system performance

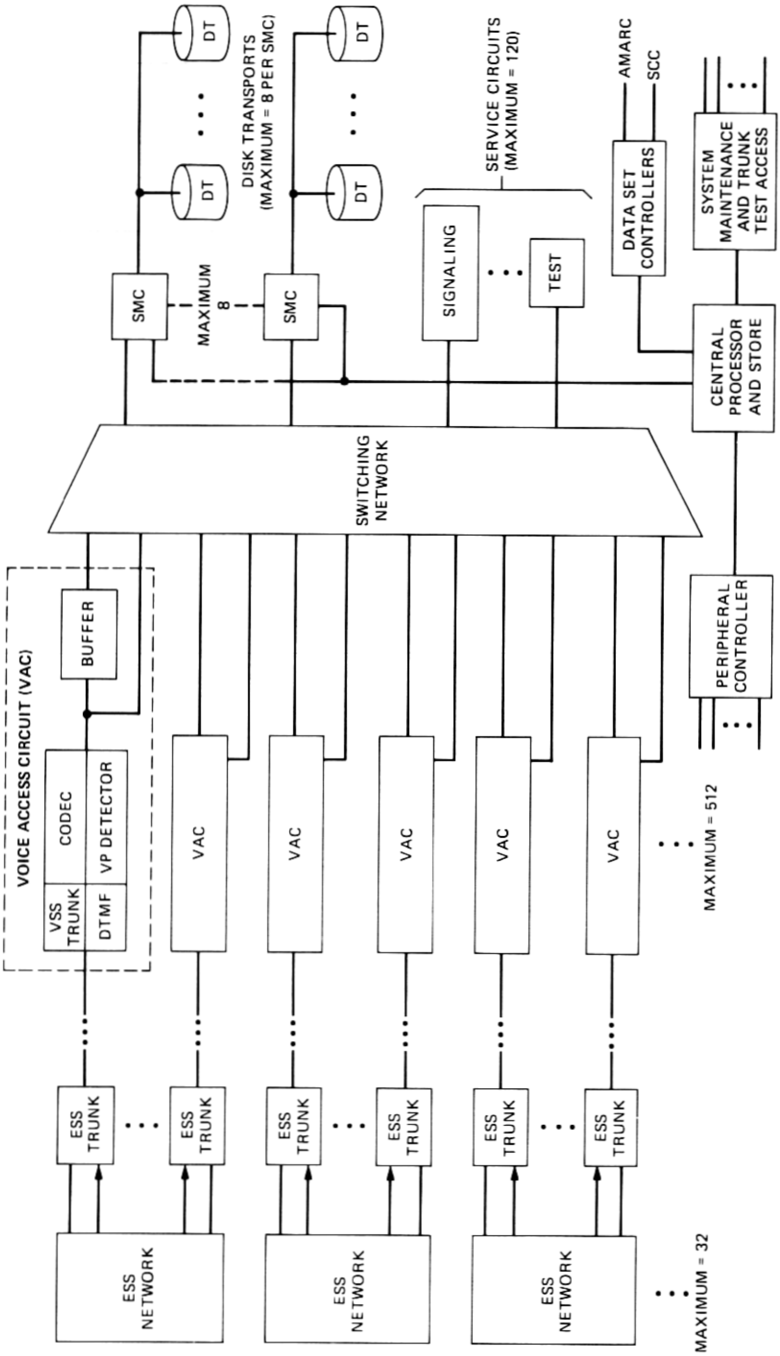

Fig. 1—No. 1A vss architecture.

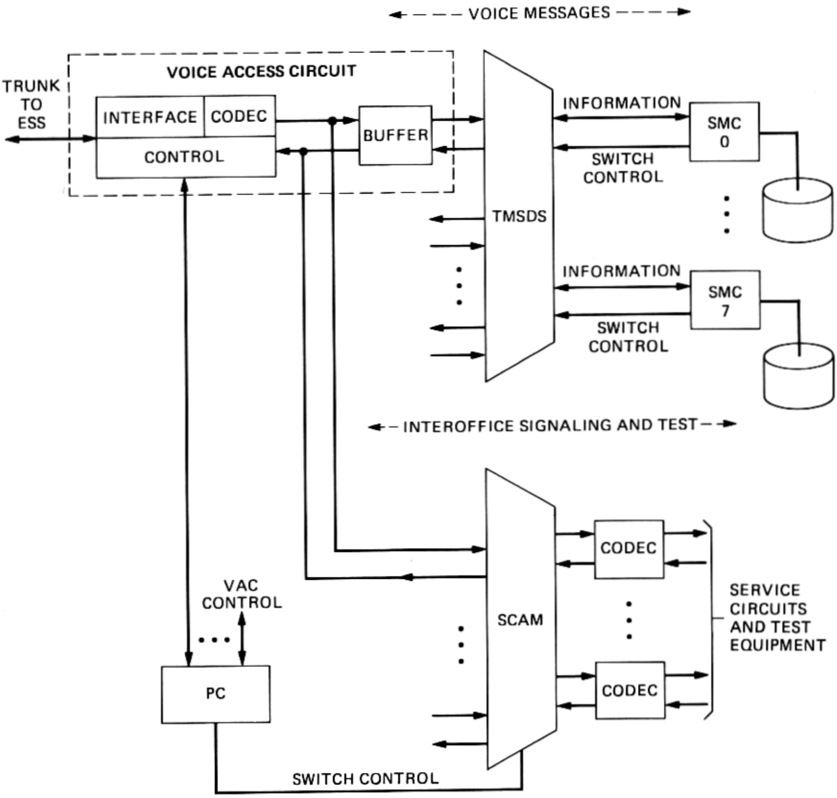

Fig. 2—No. 1A vss network partitioning.

that indicate the proper functioning of the vss office. The D schedule gives the daily totals of the C and H schedules. Also, the traffic measurements on trunks, storage, and usage are enumerated by service on the D schedule for specific half-hour periods.

2.5 *Interface to the total network data system*

The 1A vss has two independent data links which comprise the interface to the Total Network Data System (TNDS). One connection is to an Engineering and Administrative Data Acquisition System (EADAS), and the second is to the Network Administrative Center (NAC).

The EADAS polls vss for the H, C, and D schedules. Within EADAS, sets of traffic measurement outputs are triggered when preset threshold criteria are exceeded to indicate heavy traffic conditions. The EADAS also supplies the vss traffic measurements to other TNDS systems for additional traffic processing and analysis.

The NAC channel is an interface to the No. 2 Switching Control

Center (SCC) system over a dedicated facility. It is used to transmit the Q schedule. In the event of an EADAS data link or machine failure, the NAC channel serves as the backup for all traffic data (H, C, and D) in addition to the normally received Q schedule.

2.6 Transaction file

The transaction file is a mechanism for the collection, storage, and retrieval of data generated by the VSS application software. Raw data concerning the software activity or state at given times during a program's execution are written to the disk file. The transaction file is processed to produce a detailed engineering, service cost, and human factors characteristics of VSS services. Included in this information are distributions of holding times and storage utilization. The traffic measurements give the average values of the traffic parameters, but the transaction file analysis gives the complete distributions.

III. RELIABILITY

3.1 Overview

Although experience with previous ESSs may be drawn upon in establishing reliability criteria for 1A VSS, there are important differences. For 1A VSS, the term "call" has a special meaning: In a switching machine, the primary job is to establish a path (talking connection) over which two parties may converse. It does not matter whether the parties are human or computer; in either case, once a connection is made, the burden of the information exchange is the responsibility of the parties, not the switching machine. The 1A VSS, however, handles one end of the information-exchange transaction; furthermore, information flows alternately in each direction as is typical of a conversation. In the course of handling a typical call, 1A VSS receives input control signals at various times, outputs system announcements in response, and either records or retrieves customer messages. This type of communication is complex. The system is continuously involved in a dynamic interaction with the customer during the entire course of the connection, and there are many opportunities for a malfunction to disrupt the call.

At the first level of analysis, the 1A VSS may be viewed as two primary subsystems: (*i*) a large disk storage community, and (*ii*) an access system which handles the storage and retrieval of disk information. These two subsystems are quite different in nature: the disks are electromechanical moving-head storage devices, while the access system consists of program-controlled digital circuitry. Hence, the approaches taken to achieve acceptable system reliability must take these characteristics into account.

The first approach used, as a defense against occasional failure of disk transports, is to create duplicate copies of stored messages on two different disks. The second approach is to engineer the remainder of the system, via which the disks are accessed, to a failure rate sufficiently low so that it does not contribute significantly to the probability of losing a message. The manner in which this is accomplished is discussed later. The net result is that the overall system reliability is presently dominated by the failure rates of commercial disk transports. The 1A VSS is expected to provide a grade of service comparable to that of ordinary telephone service.

3.2 General design approach for reliability

Reliable operation of 1A VSS, as with other large systems, depends upon numerous design and manufacturing factors. The high probability of long-term reliable operation is built into the system at every level—from system design through hardware, firmware and software designs, carrying through quality control in the manufacturing process.

At the system level, extensive redundancy is used. Key units (e.g., the Auxiliary 3A Processor) are fully duplicated, and are provided with automatic protection switching so that a failure of either unit will still leave the system fully functional. Engineered units (e.g., SMC and disk units), whose quantity is traffic dependent, have on-line spares. The system design also encompasses power backup, communication channel redundancy, alternate external interface channels, and trunk group diversity arrangements to obviate or minimize the impact of component failures. At the hardware level,[2] self-checking arrangements are extensively used in the design. Cyclic redundancy check (CRC) circuitry monitors all message data at several points along the storage and retrieval path. Parity circuitry checks data transmissions along communication buses into and out of main memory and over disk storage and retrieval paths. An alarm system of ferrod scan points keeps watch over power and other hardware conditions. Firmware in the intelligent controllers [SMC and Peripheral Controller (PC)] routinely performs a number of checks to assure proper setup of switch and matrix connections, as well as validity of the transmitted data.

Furthermore, the analog path over the trunks which connect 1A VSS with ESS offices is automatically checked by use of an interoffice communication scheme which uses MF signalling over these same trunks.

At the software level, 1A VSS incorporates an extensive system of maintenance, diagnostic, and audit programs which can detect and locate a wide variety of hardware faults and transient errors. The maintenance system will perform appropriate system reconfigurations if necessary.

It should be noted that 1A VSS reliability, and hence availability, is influenced by the time required to repair certain faults. If one of a duplicated pair of units has failed, the system still functions normally. However, it is now vulnerable to a total outage which would occur if the mate unit should fail before the first one is repaired. The probability of this happening is proportional to the repair time. The 1A VSS design incorporates a number of features which aid craft personnel to minimize repair time. These include automatic trouble-locating circuits, firmware and software, trouble-locating reference material to interpret diagnostic results, modes of manual diagnostic program execution, built-in test facilities, and a physical arrangement in which circuit boards and 95 percent of the interunit cabling are connectorized.

The quality of the 1A VSS equipment itself is controlled by tests and inspections at each assembly level through the manufacturing and installation process. Basic components are either manufactured by Western Electric or purchased under rigid specification to guarantee the incoming quality. Once components are assembled onto circuit modules, these are thoroughly tested on computer-driven test facilities. When the modules are assembled into functional units and frames, these are, in turn, tested at the factory prior to shipment. Since a standard floor plan is used by 1A VSS, the connectorized cables which will be used in the field to interconnect frames are included in the factory test and shipped along with the frame.

Final testing occurs at the 1A VSS office site where, as part of the installation process, the frames are interconnected and operated as a system.

3.3 Hardware redundancies

3.3.1 Duplicated units

On 1A VSS (see Figs. 1 and 2), there are two unit types which perform common-control functions on which system operation depends.[2] These are the Auxiliary 3A Processor (AP) and the PC. Each is fully duplicated.

Associated with each AP is a complete set of communication buses over which it exchanges data with other units. Communication with the SMCs is via a direct-memory access (DMA) parallel channel. The even-numbered SMCs connect via one set of bus hardware, while the odd-numbered SMCs connect via the second set. This arrangement is replicated for the second AP. The *A* and *B* power buses similarly supply odd- and even-numbered SMCs. Taking this arrangement into account, the 1A VSS software always places the duplicate of a message on a disk subsystem (SMC and its disks) of the opposite group (even or odd). Therefore, even if an AP should fail, and then one of the DMA channels on the good AP should also fail, the system will be able to retrieve a copy of every stored message. Since each AP also has a separate DMA bus connection to each PC, the system can survive a failure of any PC, AP, or interconnecting parallel channel.

Another characteristic of these redundant bus arrangements is that a failed peripheral unit can always be diagnosed from the active AP. This means that the standby AP can remain at all times in the update mode; in this way, the system is better prepared to handle an AP failure as well.

Each of the PCs, which controls the per-trunk circuitry and service circuit access via the Service Circuit Access Matrix (SCAM), has its own communication channel and interface circuit in each Voice Access Circuit (VAC). Hence, failure of either a PC or its communication with a VAC can be bypassed by switching to the other PC.

3.3.2 Engineered units—M plus N sparing

Several unit types are engineered to a quantity determined by the traffic to be handled by the office. This, of course, includes the VAC units which contain the per-trunk circuitry for up to 16 trunks, and the service circuits, such as MF transmitters and receivers. The disk subsystems (SMCs and their disks) are also engineered, although the quantities can also be influenced by other factors such as reliability considerations or traffic measurements. For example, the number of disks provided will depend on the average length of time that messages are left in the system. In like manner, the Time Multiplex Space Division Switch (TMSDS), which interconnects the trunk and the disk equipment communities, is designed in a modular fashion for convenient growth.

A failure of one of these engineered units could degrade service, but it could not take the entire 1A VSS out of service; hence, the number of spares provided is less than full duplication. The sparing philosophy used has come to be known as M plus N sparing where, for M in-service units, a number of spares N is provided where $N < M$. Full duplication is where $N = M$. An additional philosophy applied for VSS is that the spares are kept in service and in use. In this arrangement, no unit can be identified as the spare; however, should a failure occur, the number remaining in service can carry the traffic at the engineered level.

There are two main advantages to keeping spares on line: (*i*) they contribute to a better grade of service during the large fraction of the time when there are no failures, and (*ii*) they are periodically tested by both routine diagnostics and by the operational checks which go on all the time in in-service equipment.

3.4 Automatic trouble detection and protection switching

As alluded to earlier, a key element in the graceful failure mechanisms of 1A VSS is the action taken by the maintenance and diagnostic software. These programs react to either hardware or operational software indications of trouble by testing suspect units or communication paths, isolating the trouble area and reconfiguring the system so as to bypass the faulty equipment. The diagnostics are also routinely called in by the maintenance software for testing of the system each night during hours when the system is relatively idle. The maintenance software is described in Section IV.

3.5 Storage duplication

Whereas 1A vss disks hold programs and data used by the processor, the bulk of the storage media is required for messages and announcements. In general, when any information is stored, it is then scheduled for duplication on another disk associated with another SMC. This arrangement allows outage of a disk or of other system components without loss of any data.

Although processor information duplication is scheduled with a high priority, analysis shows this to be unnecessary for voice messages. That is, as long as messages are duplicated within approximately one hour, the delay contributes little to the probability of lost messages since this delay is still very short relative to the mean time between failures of disk transports.

However, the advantage of the delayed-duplication philosophy, is that priority may then be given to more urgent service-affecting activity, such as the handling of calls during a temporary peak traffic period. During much of the time when the system is not heavily loaded, messages will be duplicated quickly.

IV. MAINTENANCE SOFTWARE

4.1 Overview

A complex maintenance software system is required to enable the 1A vss to meet the high reliability objectives described in Section III. The primary responsibility of this maintenance system is to accept error indicators from the operational and administrative software, reconfigure the 1A vss such that the suspected faulty unit is isolated, diagnose the isolated unit without affecting normal 1A vss functions, and help resolve exactly where the fault is in terms of replaceable circuit modules. The main interface with the operational and administrative software is the error control subsystem. This subsystem receives the error indicators and determines the corrective action. If the error control subsystem is unable to maintain a working configuration, the system recovery subsystem is called into action. Other parts that make up the maintenance software system are (*i*) The trunk maintenance subsystem, which includes the ability to perform automatic tests between the 1A vss and connecting client offices, manual tests from either end, and automatic trunk administration functions. (*ii*) The routine diagnosis subsystem, which is responsible for automatically testing the entire 1A vss periodically. (*iii*) The power/alarm subsystem, which monitors the system's power and alarm indicators, controls the system status panel, and operates all the other audible and visual indicators. (*iv*) The maintenance administrative subsystem, which accumulates hourly, daily, and monthly maintenance status and reports. All of the above functions are controlled and administered by the maintenance control subsystem. This control subsystem coordinates all maintenance activities in the 1A vss, maintains the vss configuration database, initiates all diagnostic executions, and administrates system reconfigurations. The remainder of this section will describe these subsystems in greater detail.

4.2 Error control and error history analysis

The error control subsystem receives error indicators from operational and administrative software. It is responsible for evaluating these error inputs and determining what action to take. Figure 3 shows the sequence used to maintain a working configuration when the system experiences errors. Error messages sent to error control can be classified into the following types: (*i*) A device handler has lost the ability to communicate with its hardware. (*ii*) An error occurred but

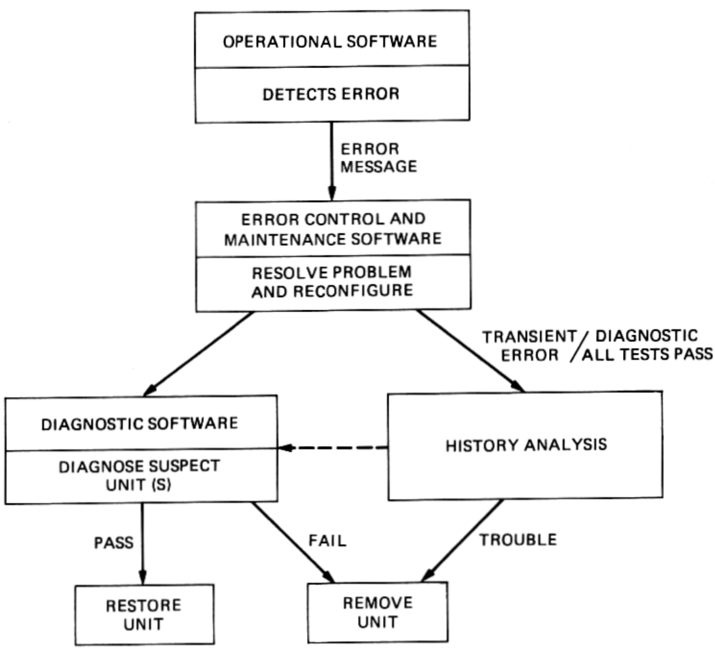

Fig. 3—Error detection/fault recovery.

the operation was successful because of a retry strategy. (*iii*) An error occurred and either no retry strategy was employed or it was unsuccessful.

When an error message of the first type is received, the unit involved is already out of service since fault recovery code is designed into the device handlers to cause an immediate reconfiguration if communications are lost with the unit. In this case, error control does not have to resolve the problem. The action taken is to request a conditional restore of the unit identified in the error message. The unit is already out of service so this request results in a diagnosis of the unit. If the resulting diagnosis fails, appropriate failure information used to resolve the fault is displayed on the maintenance teletypewriter (TTY) and the unit is left out of service. Manual action will then be required to restore the unit. If the unit passes diagnosis, it is put back in service and the history analysis part of error control is informed. In certain instances, it is possible for a unit to fail but pass diagnosis and be restored. To prevent an endless cycle of this type of occurrence, each time the unit is restored, history analysis is informed. If this cycle repeatedly occurs in a short time span, history analysis will order the unit to be removed from service and to be put in a trouble state. Manual action will then be required to restore the unit to service.

An error message of the second type is called a transient error. No attempt to resolve the problem is made by error control. Instead, the transient error is sent to history analysis. If many transient errors for a particular unit occur in a short time span, history analysis will cause the unit to be diagnosed by requesting a conditional restoral. Since the unit is in service at this time, it will be removed from service when the restore request is received (diagnostics can be run only on an out-of-service unit). Once the unit has been removed from service, the actions taken are identical to those in the first error-type sequence, including the feedback to history analysis if the diagnostic passes.

Error messages of the third type trigger the greatest amount of activity. Some of these messages are such that a specific unit is immediately implicated. For these cases, a conditional restore of the unit is requested. Other type-three error messages do not allow error control to make an exact determination of where the fault lies. These messages result in conditional restorals of some suspect units, and implication of other units via transient error messages sent to history analysis. In the former case, the resulting diagnostic will either find a fault or exonerate the units. The action in the latter case is identical to that taken for transient error messages sent by operational software.

It is expected, in the cases where the determination is not clear, that the unit implication lists and history analysis thresholds will be optimized as field experience becomes available.

4.3 System recovery

The system recovery subsystem has the responsibility for recovering a working 1A vss when the system must be reinitialized (similar to phasing the ESS machines) because of processor problems, when catastrophic peripheral hardware failures occur, or when 1A vss software insanity occurs. When the system is reinitializing because of processor problems, system recovery has little control over the situation. If sufficient time does not elapse between successive reinitializations, the reinitialization level escalates. At each level, more drastic software initialization occurs until the bootstrap level is reached. Once this occurs, system recovery takes its first positive action—it triggers a memory reload of all resident memory programs and the office database.

As specified in Section III, some of the 1A vss periphery is duplicated. If a duplex failure occurs, the failure is catastrophic and 1A vss ceases to function operationally. When the duplex device handlers recognize this occurrence, they send a message to system recovery and stop functioning. System recovery will attempt to find a working combination of one of these devices and one of the processors by trying all possible combinations and by driving the initialization level higher and higher each time. Unlike the escalations caused by processor problems, if the system recovery subsystem is unsuccessful at finding a working combination, the escalations can be stopped manually by pressing the manual stop key on the system status panel. When the manual stop key is pressed, the 1A vss will stop thrashing through initializations and will settle into a quiescent nonoperational state where diagnostics can be run manually on the duplex-failed unit until the faulty component is found and repaired. When one of the units is functional, the manual stop key is released and a manual bootstrap will restart the system.

The system recovery subsystem provides an outlet for software modules (including the system recovery modules) that find themselves in untenable states from which further processing would cause further system software failures.

When this occurs, a module can send a message to system recovery which causes an initialization of the module, and perhaps all other modules depending on the level of initialization. This could cause an escalation to the memory reload bootstrap level. The manual stop key has no effect on software sanity initializations.

The last interface to system recovery is utilized when the 1A vss configuration has deteriorated below a predefined threshold. For threshold analysis, the 1A vss is divided into the storage subsystem, the service circuits subsystem, and the voice access (per trunk) circuit subsystem. When 50 percent or less of the units in any of these categories are in service, system recovery is notified. The first notification will cause a switch of the processors which will raise the level of initialization by one. The next unit removed from service in that group will trigger conditional restores of all the out-of-service units in the affected subsystem. Further reports of units being removed from service in this category will be ignored until all the units in the category are out of service. If this occurs, all the units in the category will be unconditionally restored to service (unconditionally means that no diagnostic is run), and the threshold recovery algorithm is reset. The algorithm is also reset if sufficient time elapses between any of the levels of recovery. The manual stop key functions as described in the duplex failure case.

4.4 Trunk maintenance

The 1A VSS office connects via trunks to the 1/1A ESS connecting offices. The main philosophy of the 1A VSS trunk maintenance plan is that the connecting offices, regardless of type, are the controlling offices for the interconnecting trunks. This means that the 1A VSS will have supporting automatic trunk test equipment, but will not have responsibility for initiating trunk facility maintenance. The connecting offices execute and evaluate end-to-end operational and transmission tests. Nevertheless, the 1A VSS does contain a substantial amount of trunk maintenance software which is used for trouble detection, trouble verification, sectionalization, repair, repair verification, service protection, and new circuit installation or circuit rearrangement testing.

The 1A VSS will accommodate the execution of end-to-end operational tests initiated from the connecting office. This test, which is run whenever the trunk diagnostic is executed at the connecting office, will validate the interoffice signaling capability and continuity over the transmission path, but will not test transmission quality. This will occur automatically at least on a daily basis. The 1A VSS will also accommodate a Remote Office Test Line/Centralized Automatic Reporting on Trunks (ROTL/CAROT) automatic transmission test. Only terminating test equipment exists at the 1A VSS; therefore, the transmission test must be initiated at the ESS connecting office end either manually or automatically through a CAROT facility. This test, which includes a verification of transmission quality, will be run automatically at least once a day.

The 1A VSS provides extensive trunk service protection on its own. Whenever an internal 1A VSS problem exists which affects one or more VACs, the corresponding trunks at the VSS end are put in the reverse-make-busy state. A trunk is in this state when it is seized from one end, but no signaling is sent or accepted by that end. Viewed from the opposite (ESS) end, the trunk is said to be high and wet. This will keep the connecting offices from using the affected trunks at the price of making them execute some software which deals with the high-and-wet state. When the internal 1A VSS problem is cleared, the trunk circuit will be put back into a normal idle state.

All carrier groups connected to a 1A VSS have hardware Carrier Group Alarms (CGAS). A CGA is reported to a trunk maintenance module which immediately removes all associated trunks from service and releases any service circuits tied to these trunks. This protection is required because a 1A VSS has only a small number of service circuits which could all be occupied by a faulty carrier that causes all the associated trunks to be seized. When the CGA is cleared, all the associated trunks will be put back in service automatically.

Whenever a VAC is put back in service, an operational end-to-end test is run on the corresponding trunk from the 1A VSS end. This test does not require any special software at the ESS end. It is used to determine if the trunk is high and wet (seized permanently from the ESS end), or if it is otherwise faulty so that it should be locked out at the 1A VSS end. The locked-out condition occurs if the ESS does not respond to a trunk seizure from the VSS end. A locked-out trunk will not be used operationally by the 1A VSS but will be treated normally if seized by the connecting office. Trunks in the locked-out state will be automatically and periodically tested by the 1A VSS end-to-end test. If it passes, the trunk will be restored to a normal condition. A trunk found to be high and wet will be restored to service automatically when the permanent seizure from the connecting end is dropped.

The trunk maintenance subsystem also provides a manual trunk maintenance capability from either end of the trunk. The craft personnel at the connecting end can request several test signals from the 1A VSS, or can be connected to the trunk test panel at the 1A VSS. By using various TTY input messages and the trunk test panel, craft personnel at 1A VSS can call test signal generators or the trunk test panel at the connecting office end.

The key points of the trunk maintenance plan for the 1A VSS are that the extensive trunk maintenance software, along with the error control/history analysis software described earlier, provide a powerful sectionalization tool for all trunk-related problems. Consequently, almost no end problems will require manual restorals of trunks at the end where the problem did not exist.

4.5 Maintenance control

The maintenance control subsystem orchestrates all maintenance activity in the 1A VSS periphery. The processor contains its own maintenance system. The maintenance control subsystem gets requests to remove units from service, to diagnose units, restore units to service, and to switch duplex units. These requests come from the error control, system recovery, trunk maintenance and routine diagnosis subsystems, and from manual inputs via the TTY from craft personnel. Maintenance control is responsible for prioritizing these requests and ensuring that no interference occurs between concurrent maintenance activities.

Requests to remove units from service will be evaluated to determine the effect of this action. Removing a unit from service may affect several other units, or it could trigger particular alarm conditions or even system recovery action. Maintenance control must make these determinations and take appropriate action. One example of an interactive condition is the removal of an SMC from service, which will necessitate the removal of all disks connected to it, along with its associated second-stage switch (one of the modules in the two-stage TMSDS shown in Fig. 2). All three of these unit types are grouped into a family, and anything that affects one member of the family is evaluated to determine its affect on other family members. There are several other family groupings in the 1A VSS.

Another situation arises when the removal from service of an SMC would leave the system below a critical minimum number of operational SMCs (as detailed in the system recovery section). In this case, a routine remove request would be denied.

Requests to diagnose units will cause these units to be removed from service if they are active, and then helper units will be selected as required before the request is passed along to the diagnostic control subsystem. Requests to restore a unit to service will first cause it to be removed from service if required, and then diagnosed. If the unit passes the diagnostic tests, it will be initialized and restored. A unit can also be unconditionally restored which will cause it to be initialized and put back into service without being diagnosed. When a unit is restored, a family evaluation will take place. For example, this occurs when all the disks connected to a particular voice message controller are taken out of service; then the storage media controller and its associated second stage switch are also taken out of service but are put in a nonfault stage. When any one of the disks is restored to service, the associated SMC and second-stage switch are also restored automatically.

4.6 Diagnostics

The 1A VSS peripheral diagnostic programs are used to detect faults in their respective units; the resulting failure data are then used to identify any of the replaceable circuit modules which are faulty. A table-driven diagnostic design is used, and a high-level-macro approach facilitates the diagnostic development. The same diagnostics are used for frame testing in the factory, for installation testing at the field site prior to cutover, and for on-line testing programs, while the system is operational. The diagnostics are triggered either manually, by craft personnel using a TTY, or automatically as was discussed earlier. The manually initiated diagnostics may be originated from either the on-site maintenance TTY or remotely from a Switching Control Center (SCC). The trouble location capability is an integral part of the diagnostics and will be described later. In general, the location capability is at its maximum when a unit contains a single fault; if it contains multiple faults, the location resolution is reduced.

Each unit diagnostic is a collection of macros used to test the unit. These macros expand into data table words when assembled. Each macro expands into an OP-CODE or INDEX word, plus one or more data words. These data table words are grouped into one or more segments each containing less than 2K words. When a diagnostic is triggered automatically, all the segments in a particular unit diagnostic are run on the unit under test. If the diagnostic is manually triggered, all segments or selected segments can be run. A diagnostic segment can also run in an interactive mode where execution proceeds to a selected point in a segment, or proceeds from a selected point to another, or loops over one or more tests.

Figure 4 shows the structure of the diagnostics. The data tables reside in auxiliary memory in the 1A VSS disks. When executed, the segments are overlayed, one segment at a time, into a 2K paging buffer in main memory and are used to drive the diagnostic control program which resides in main memory. A simple model of this control program is included in Fig. 4. After overlaying a segment of data tables, a task dispenser examines the first data table word and extracts the OP-CODE or INDEX information. This is used to pass control to the appropriate task routine. Each macro type has a corresponding task routine which fetches the data words following the OP-CODE or INDEX word and executes accordingly. Most of the task routines build commands for the 1A VSS peripheral devices and then send these commands to the appropriate device handlers. These handlers send the commands to the devices and they in turn cause the execution of one or more device-resident firmware routines. The result of this execution is sent back to the task routine via the device handler. If this sequence was triggered by a test-type macro, the task routine will compare the results with an expected result (supplied in the data words of the macro expansion) and will determine whether the test passed or failed.

Of special interest is the fact that the diagnostic structure resides in three different memory media. The data tables are in auxiliary mem-

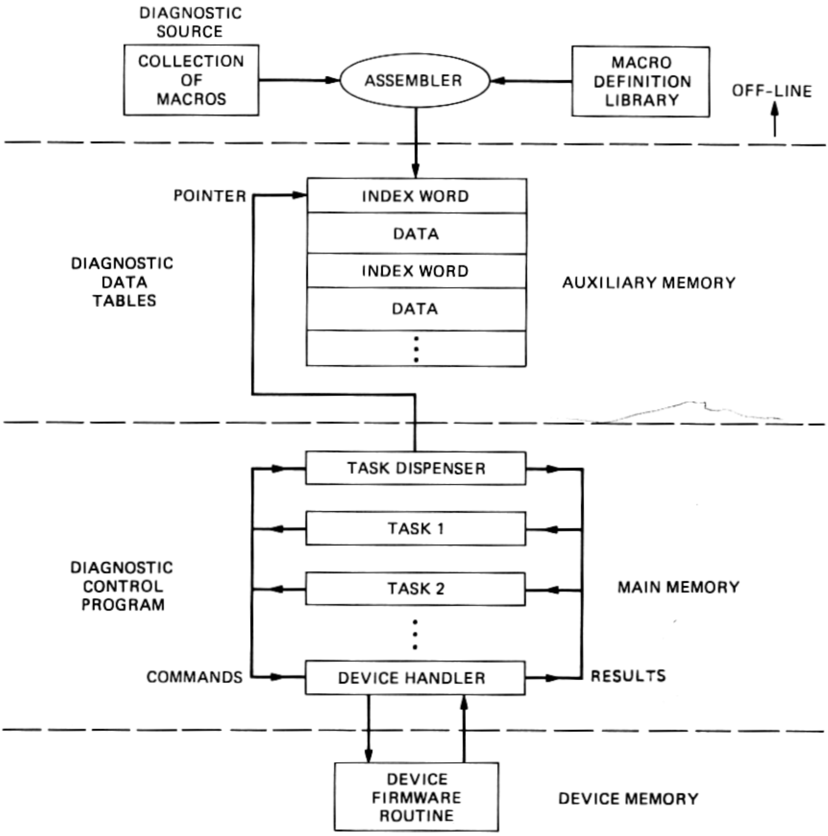

Fig. 4—Diagnostic design.

ory, the control is in main memory, and the work routines are in device memory (firmware). This table-driven, high-level-macro approach to designing diagnostics has been proven very successful in other ESS developments.[3] It has permitted the diagnosticians to concentrate efficiently on the functional requirements of the hardware.

4.7 Trouble location

Trouble location uses all available information to deduce where a fault exists in the system. Once the resolution is made, the fault is corrected and the out-of-service unit is restored to service. Trouble location may require analyzing error conditions, plant measurements, maintenance reports, and any other indications of atypical system behavior. No attempt will be made here to describe all possible approaches. However, the great majority of all faults will be resolved by the unit diagnostics.

The trouble location information is built into the diagnostic source.

A trouble location macro follows each test-type macro in a diagnostic segment. This macro requires the diagnostician to comment on what the test was checking for, to list the suspected replaceable circuit modules in a prioritized order, and to specify any special instructions required for the repair process. When the diagnostic source is assembled into data tables, these macros are expanded into readable repair information in the diagnostic listing. The listings are processed off-line where the trouble location information is extracted and used to form a separate Trouble Location Manual (TLM). When a diagnostic detects a test failure, it terminates at that point. Craft personnel can use the test number and segment number to find the appropriate diagnostic listing or the TLM where the trouble location information will be found.

4.8 Routine diagnosis

Periodically, the routine diagnosis subsystem will test the 1A VSS in an attempt to provide automatic preventive maintenance. Periodic testing will detect a fault soon after it occurs, which reduces the probability of the occurrence of a second fault before the first one is repaired. This is desirable since, as previously mentioned, the power of the diagnostics is reduced if the unit under test contains multiple faults. Daily routine diagnosis will attempt to remove, diagnose, and restore all the 1A VSS units that are in service. If the diagnostic fails, the unit will be left out of service.

There are several special considerations that the routine diagnosis subsystem must observe. It will run only during the intervals when the system has little or no customer load. It will cause the processors to switch only after first triggering a diagnosis of the processor that was the standby. It must take hardware family associations into account to minimize the effect on the rest of the system.

Routine diagnosis also has the responsibility for performing certain test sequences that are not executed when the diagnostics are run as previously described. These test sequences can be run manually if the appropriate diagnostic parameters are used, but they seldom will be executed this way because of their long execution times. Two examples of such sequences are the multiple path select tests in the switch and matrix diagnostics, and the complete media test and initialization portion of the disk diagnostic. These sequences will be run periodically, but not on a daily basis.

4.9 Power/alarm control

The power/alarm subsystem monitors the system's power and alarm indicators (with the exception of the carrier group alarms described in Section 4.4). Examples of these are unit power indicators, bus power alarms, and building alarms. The power/alarm software also provides routines for other maintenance software to trigger minor, major, and critical alarms in the system. These power and alarm inputs will result in appropriate TTY output messages, audible and visual indications, and corresponding indications on the System Status Panel (SSP). All of these alarms will be sent to the connecting SCC system using another port on the maintenance TTY channel and telemetry. An interesting feature of the 1A VSS maintenance philosophy is that the power and alarm indicators which provide input to the power/alarm subsystem will not trigger any maintenance activity other than the indicators described above. The reasoning behind this is that all such power/alarm problems will result in operational problems with the unit or units involved and will be handled by the error control subsystem.

Most 1A VSS peripheral units have power and status indicators built into the power switch module visible on the unit's front panel. These indicators show if the unit is operating normally, is out of service, is powered down or is in a power alarm state. The indicators are controlled by the power/alarm subsystem. The status of the 1A VSS units is also displayed on the SSP. If all the members of a particular unit type (e.g., SMCs) are in service, the corresponding light on the SSP is extinguished. If one or more units are out of service, the light is lit. Minor, major, or critical alarms result in output messages describing the problem, appropriate indicators on the SSP, audible alarms and aisle pilot lights indicating which aisle or equipment is experiencing the problem. If the 1A VSS does not have resident craft personnel, the audible alarms can be shut down at the site and monitored only via remote connections to the SCC system.

4.10 Administrative maintenance function

A great deal of effort was put into the man-machine-interface design of the administrative maintenance capability. This capability consists of:

(*i*) Reporting features which allow the craft personnel to ask for the maintenance status of various system components via TTY input messages.

(*ii*) Automatic outputs of maintenance information on an hourly basis.

(*iii*) Automatic outputs of detailed plant measurement data on a daily basis.

(*iv*) Automatic outputs of a summary of the plant measurement information on a daily basis (used by the management in charge of the system).

(*v*) Monthly measurements and summary records.

The hourly maintenance information can be tailored by the individual system managers. The default case is to report all available information. The entire report or selected categories can be printed out or inhibited as desired. Examples of items comprising this report are a list of units out of service, entire trunk groups that are inoperative, and carrier facilities in the alarm state. The detailed plant measurements show which units had transient error conditions, which were automatically removed from service, which were found faulty as a result of the automatic removal, and the measurements also show the length of time each remained out of service. Other plant measurements indicate the frequency of system reinitializations and the levels at which the reinitializations occurred, and provide a calculation of the message reliability based on the Mean Time to Repair (MTTR) and the Mean Time Between Failure (MTBF) of the storage subsystem.

V. INSTALLATION TESTING AND ACCEPTANCE TESTING

5.1 Overview

As the final phase of the installation, a new VSS is put through an extensive series of tests to verify proper operation. Most of these tests result in TTY printouts of system actions that can serve as a permanent record of test results. In addition, a high-temperature test may be run at the option of the BOC. The BOC personnel either participate in these tests or subsequently review the results. At this point, the system is turned over to the BOC for final acceptance testing prior to cutting the system into service.

Acceptance tests are run by a BOC to assure itself of the proper operation of a newly purchased system. It is highly likely that no two VSSs will ever be exactly alike in terms of installed office configuration, trunking arrangements, connecting office number and types, and the mix of customers served by these offices. Nevertheless, the BOC has extensive standard test documentation and the expertise of the Western Electric installation organization to draw upon in tailoring testing to the specific VSS office configuration.

Acceptance tests may include:

(*i*) an audit of the installation tests,

(*ii*) tests which check the interaction of the 1A VSS with other systems to which it connects (e.g., ESS client offices and operational support systems).

The installation tests are described in Section 5.3. In these tests, the system's ability to handle a heavy calling volume is verified by the application of a simulated traffic load. The equipment which generates these test calls is described next.

5.2 Call simulation equipment

The equipment used to provide a simulated traffic load to 1A VSS is different than that used in ESS offices because the requirements of 1A VSS are unique. The processor on either an Electronic Switching

System (ESS) or on 1A VSS can be presented a load by generating a large number of short-holding-time calls with computer-driven test facilities. In the ESS case, lengthening the holding time on calls would additionally load down the network. This is not generally done—it is considered unnecessary because networks are well understood and can be engineered to a highly predictable level of performance. Long calls would, incidentally, also load down the test facility; a much larger computer would be required to manage the much greater number of calls concurrently extant within the switching system.

For 1A VSS, both a high calling rate and long call holding times are required. It is true that the processor, which handles the call setup and tear-down, is loaded by a high calling rate. However, unlike the situation in an ESS, the 1A VSS processor also has work to do during the course of a call to handle the storage and retrieval of messages. Furthermore, the SMCs and disks are loaded most heavily by lengthy messages which must be stored and retrieved.

The corresponding throughput capacity requirement on the call simulation equipment is large. This is handled by a number of microprocessor-driven call simulators which provide a distributed load. In each VAC unit (of which there are two in each VAC frame) there is a normally unoccupied circuit module location having backplane wires to each Trunk Access Circuit (TAC). A call-simulator board may be plugged into this location to provide a call load on any or all of the 16-trunk circuits in this VAC. Simulated traffic is applied to selected trunks by the insertion of a strap plug (in lieu of an actual trunk connection) for each such trunk on the VAC backplane. Each call simulator can generate traffic autonomously, or can be monitored and controlled from a common facility.

The call simulators are capable of generating any or all of five call types associated with the Call Answering service.[1] They produce messages comprised of a pulsating tone that is encoded to contain a check number; message length may be set from 10 to 80 seconds in 10-second increments. Front panel switches (see Fig. 5) or external control may select the call type and message length parameters. Additionally, the E&M lead signalling is handled as is the generation of multifrequency (MF) digit strings which contain call type and directory number information for the VSS. Retrieved messages are decoded and matched for a correct check number. Counts of total calls and of errors in each of five categories are recorded by the microprocessor. They are also selectively displayed on a three (hexidecimal) digit read-out on the front panel.

At the shortest (10-second) message length, a call simulator is capable of generating in the vicinity of 160 calls/hour/trunk depending, of course, on the response time of the VSS. This substantially exceeds

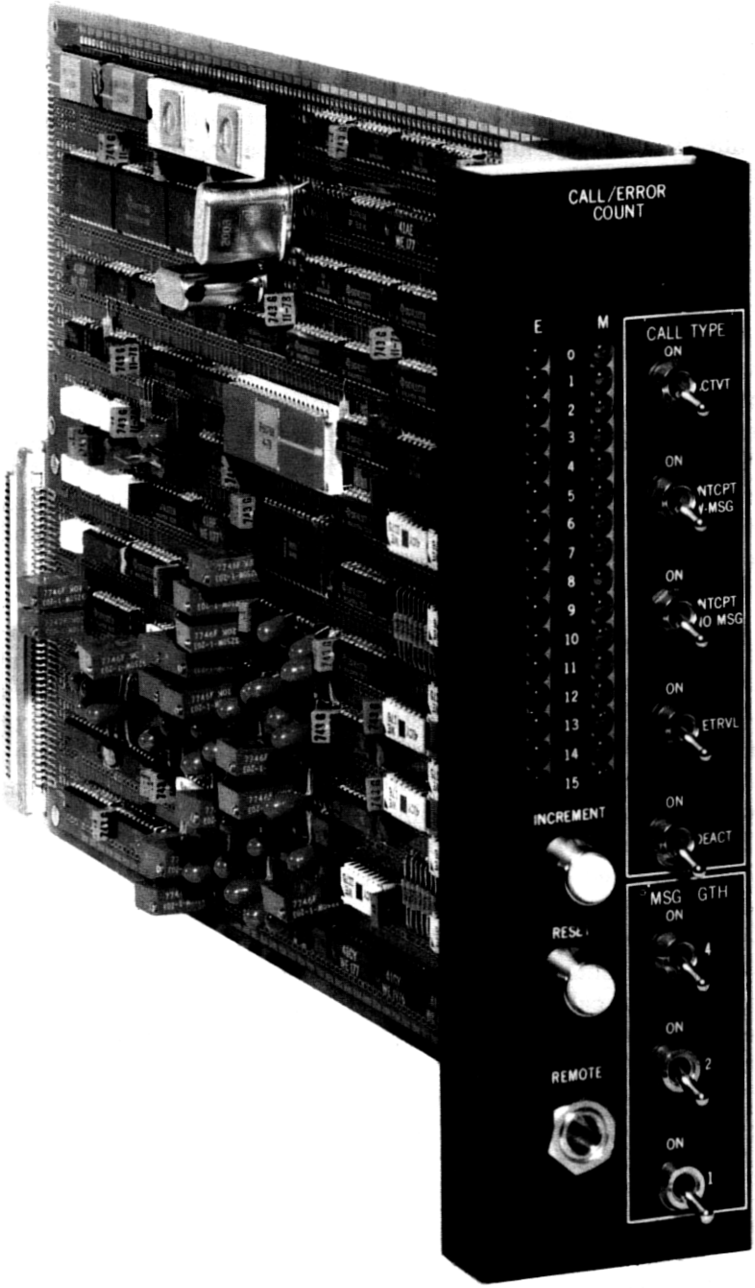

Fig. 5—Call simulator.

anticipated per-trunk calling rates which are predicated on longer call holding times. Therefore, in doing volume tests, message lengths are chosen appropriately to obtain the calling rate for which an office has been engineered.

A jack on the front panel provides a 300-baud RS 232 communication interface to a common monitor/control unit. This unit contains a minicomputer controller and a teletype for output. A maximum of 32 call simulators, each located in a VAC unit, would be concurrently connected to this monitor/control unit in a maximum-size 1A VSS. In the normal monitor mode, a full report is issued every 15 minutes. Others modes are available for hourly, daily, or on-demand reports. For each active trunk in the system, a printed line is produced that contains counts of total calls and of each of the five error types. Control and monitoring of all call simulators from this central point greatly facilitates the administration and recording of the integrated and maintenance volume tests, as discussed in the next section.

5.3 Installation tests

Various final checks performed by Western Electric installation are discussed in the following:

(*i*) Idle System Tests are those in which the system is caused to sequence through the routine diagnostics which are normally run once a day. No traffic load is applied. This test verifies that all units and subsystems are operating properly.

(*ii*) Integrated Volume Tests involve a simulated traffic load being applied to the system. In this test, the 1A VSS is required to carry the level of traffic (calls per hour) for which it was engineered, for a 24-hour period, during which full grade-of-service requirements must be met.

(*iii*) Maintenance Volume Tests verify the system's ability to recover properly from hardware troubles and software initializations. Simulated traffic is applied during the test. The grade of service delivered by the 1A VSS during induced disturbances is monitored by the call simulation equipment and printed out on an associated TTY. The printout provides per-trunk counts of the number of calls handled properly, and of call-handling errors in different categories. The number of such errors, if any, which are allowed depends upon the severity of the disruption induced in the system. For example, the manually initiated removal from and restoral to service of any units in the 1A VSS should be accomplished without perturbation of call-handling activity. On the other hand, the disabling of an active controller will abort a call being set up at that instant, while a major software reinitialization could suspend call handling for a matter of seconds. A failure of commercial ac power is induced to verify the system's ability to transfer smoothly to battery-backup power. In general, the system design is such that the voltages supplied to VSS operational equipment are not disturbed by such a failure.

Also, various additional tests are run to cover interfaces or test equipment not checked by diagnostics (e.g., trunk test facilities).

5.4 Acceptance tests

During the interval between turnover and cutover of a system, the BOC performs acceptance tests to verify that the system will perform as engineered. Aside from the optional repetition of any of the installation checks, these tests deal primarily with interfaces to other systems. For VSS this includes the connection to the automatic message accounting recording center (AMARC) system, which receives billing information; to EADAS (discussed in Section II); which collects traffic data; and to the SCC for maintenance monitoring. However, the largest interface is made up of the trunk groups, which connect to ESS client offices.

A complete test of VSS features is made at the option of the BOC. Since the presently planned services and features can be thoroughly tested over the interconnecting trunk groups by several hundred phone calls, the development of special field test facilities has not been warranted. These intersystem tests are spelled out in a "script" which defines the specific actions (and their timing) to be taken in placing each call. A grouping of test telephones has been defined, each of which is provided with specific ESS features (e.g., one- or two-digit speed calling, call forwarding, call waiting, etc.). Appropriate calls placed on these phones can then verify that the combination of the ESS with its software, and the 1A VSS with its software, are interacting correctly so as to produce the expected announcements or other responses.

VI. SUMMARY

A conventional telephone switching office mediates communication between two parties by setting up a connection over which they may converse. Unlike a switching office, the 1A Voice Storage System (1A VSS) takes on the role of one of the conversing parties. In the course of handling a typical call, 1A VSS receives input control signals at various times, outputs system announcements in response, and either records or retrieves customer messages. Since this type of communication is complex, offering many opportunities for a system malfunction to disrupt a call, the 1A VSS design incorporates a large number of failure-defense strategies.

The quality of service is of utmost importance. This article has presented the underlying philosophies and the approaches used to obtain this quality. Several contributing factors which have been covered are adequate office engineering to assure sufficient equipment to handle peak traffic, high intrinsic reliability based upon system redundancies and defensive software strategies, and an effective maintenance software system to minimize the effect of failures on system service. The incorporation of these system features has provided a robust 1A vss system which is flexible and reliable.

VII. ACKNOWLEDGMENTS

The work described in this article could not have been accomplished without the combined efforts of numerous system designers, as well as the contributions of the Western Electric manufacturing and installation organizations.

REFERENCES

1. D. P. Worrall, "1A Voice Storage System: New Custom Calling Services," B.S.T.J., this issue.
2. R. G. Cornell and J. V. Smith, "1A Voice Storage System: VSS Hardware Architecture," B.S.T.J., this issue.
3. P. W. Bowman et al., "1A Processor: Maintenance Software," B.S.T.J., *56,* No. 2 (February 1977), pp. 264–71.