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

1A Voice Storage System:

Architecture and Physical Design

By R. G. CORNELL and J. V. SMITH

(Manuscript received July 31, 1979)

To provide new basic capabilities associated with the Custom Calling Services II (CCS II) feature package, a new system has been developed. This system, the 1A Voice Storage System (1A VSS), is a generalized resource with the ability to receive, store, administer, compose, and deliver voice messages or announcements. The capacity of this system is such that it can provide CCS II services on hundreds of simultaneous calls and serve tens of thousands of concurrent subscribers. In this article, the circuit architecture and physical design of the 1A VSS is described.

I. INTRODUCTION

1.1 Background

Custom Calling Services now available to electronic switching system (ESS) customers have been implemented through software features and the existing switching periphery associated with the ESS system itself. However, the new Custom Calling Services II (CCS II) feature package for No. 1/1A ESS customers has created the need for a new functional capability not available in the existing plant: high-capacity, high-availability, and rapid-access voice storage. Additionally, the complexity and real-time processing demands of CCS II creates the need for auxiliary processor support for the implementation of these services. In response to these needs, a new node in the Stored Program Control (SPC) network has been developed: the 1A Voice Storage System (1A VSS). The 1A VSS is a generalized resource with the ability to receive, store, administer, compose, and deliver voice messages or announcements. The capacity of the 1A VSS is such that it may be deployed in a centralized fashion, each system serving tens of ESS

Central Offices, hundreds of simultaneous calls, tens of thousands of concurrent subscribers to CCS II features, and a total community (via subtending ESS offices) of hundreds of thousands of lines.

II. SYSTEM ARCHITECTURE

2.1 Basic structural concepts

The basic attributes of the 1A VSS are that it is a self-contained stored-program-controlled entity in the SPC network, that it interconnects with subtending No. 1 ESSs via trunks, and that it provides random access from these trunks to a massive voice storage medium. Special new services utilizing these resources are implemented by SPC within the 1A VSS. It follows that 1A VSS embodies many of the structural and operational attributes of an ESS. The basic components of the 1A VSS architecture (Fig. 1) are a switching network, highly specialized trunk circuits called voice access circuits (VACs), an ensemble of high-capacity storage devices and their controllers, a central processor and program/data store, a peripheral controller, and an ensemble of service circuits and data set controllers.

Although 1A VSS interfaces to No. 1 ESS via analog trunks, all internal switching and storage of voice is implemented digitally. To support this, each VAC is equipped with a Coder/Decoder CODEC which transforms analog voice to/from a digital representation. The digital bitstreams associated with encoded voice are switched via the digital network to storage media controllers, and ultimately to the media itself. Voice storage is implemented digitally via high-capacity moving-head disk transports. Moving-head disk transports provide high-speed random access to stored messages such that the delays experienced by humans after requesting message playback are barely perceivable.

A major objective associated with the 1A VSS development has been to minimally impact the design of the ESS offices to which the 1A VSS interconnects. The 1A VSS and each subtending ESS must intercommunicate frequently to coordinate their efforts toward the implementation of service. Rather than defining the need for a new data link to be established in the ESS to effect this communication, interoffice signaling is implemented via multifrequency (MF) tone packets. The switching network (Fig. 1) is used to establish a path between the relevant trunk and one of the 1A VSS service circuits. Once the interoffice signaling phase of a call is complete, the network is used to reswitch the trunk to an appropriate storage controller.

Because the customer may control [via dial pulse or dual tone multifrequency (DTMF) signaling] the operation of 1A VSS at any time during a call, a DTMF/dial pulse receiver is permanently assigned to each VAC. This implies a tremendous scanning and control load associated with

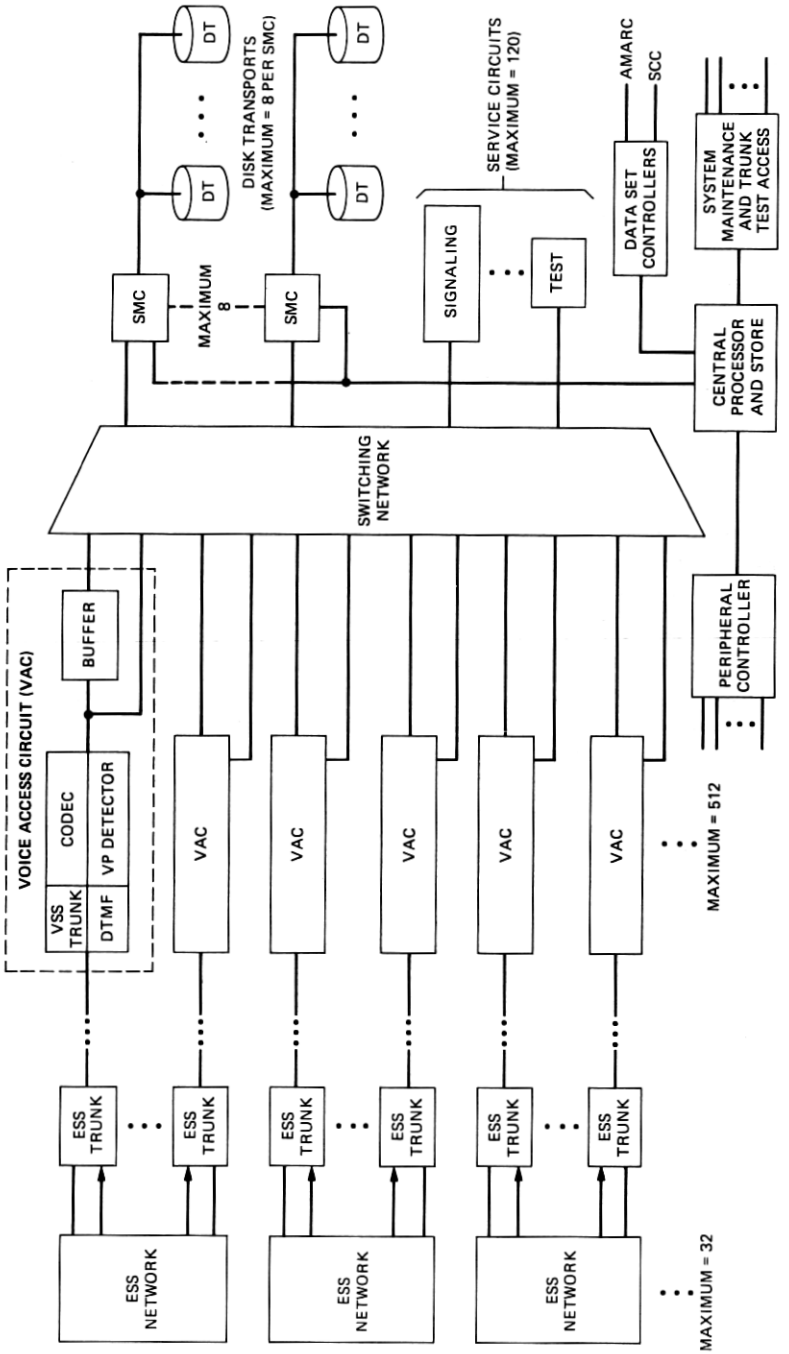

Fig. 1—No. 1A vss architecture.

the 1A VSS periphery. Much of the burden of peripheral control has been implemented via microprocessors distributed throughout the periphery itself. Each storage controller is microprocessor controlled, as is the peripheral controller. Each of these processors communicates with the central control via a functional command language which minimizes both the frequency of communication and the amount of central control real time devoted to peripheral control. This provides the central control with the additional real time required to implement the complex CCS II services.

The 1A VSS exists within the SPC network as a self-contained office. As such it has been designed to meet the reliability standards established for ESS offices. These standards include reliability and maintainability features in both hardware and software that ultimately limit system downtime to mere minutes per year per system. Specific hardware features include duplication of critical units, each implemented with self-check logic design to enable fault detection. The SPC of 1A VSS is the Auxiliary 3A Processor which was originally developed for the No. 3 ESS and No. 2B ESS systems. All signal paths experience an automatic loop-around test immediately before they are switched into a voice path. All programs and data stored on processor memory or disk are duplicated. Voice messages are also duplicated.

At the system level, 1A VSS incorporates data links to the Switching Control Center (SCC) for remote fault monitoring and control, to the automatic message accounting recording center (AMARC) for automatic message billing features, and to Engineering and Administrative Data Acquisition System (EADAS) for traffic analysis purposes. The 1A VSS trunks are automatically diagnosed via the centralized automatic reporting on trunks/remote-office test line (CAROT/ROTL) system.

2.2 The voice access circuit

Each trunk appearance at the 1A VSS terminates a VAC. The VAC (Fig. 2) consists of a trunk circuit, a DTMF and dial-pulse receiver, a voice-presence detector, an automatic gain control (AGC) circuit, an analog to/from digital CODEC, and a high-capacity, two-port FIFO buffer. Each function of the VAC is under the control and surveillance of the central processor via a microprocessor-based peripheral controller (PC). The VACs and PC communicate via a synchronous, bit-serial message protocol through transmission over an internal distribution network.

The 1A VSS VACs utilize standard Type II E&M trunk circuitry and are available in 2- and 4-wire versions.

An AGC circuit is included to provide for constant playback sound level for both prerecorded announcements and prompts and for messages recorded directly from customers. The AGC circuit also meets a

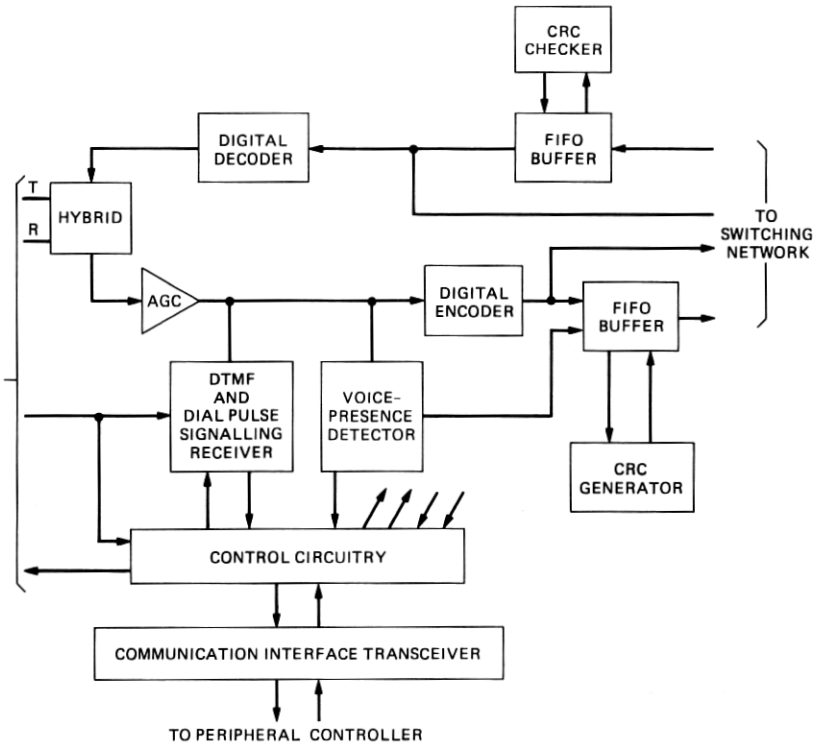

Fig. 2—No. 1A vss voice access circuit.

transmission requirement that all messages be played back at −20 VU (measured at the VAC trunk port) to assure that customers hear 1A vss playback at a pleasing level, regardless of variations in level at the time of message recording. The AGC circuit is under program control in that its function may be defeated at the command of the PC. In the disabled mode, the AGC circuit provides a fixed gain, enabling trunk maintenance testing.

At the time of the message's recording, the speaker may leave periods of silence preceding and following his message. These periods of leading and trailing silence are detected and removed by 1A vss to provide for a pleasing and efficient playback. The voice-presence detector in the VAC seeks to distinguish between signals with the characteristics of voice and signals with the characteristics of noise. Additionally, the voice-presence detector reports "voice present" for continuous signals at high levels to allow for the storage of frequency shift keying (FSK) encoded data. The voice-presence detector functions by observing the rate of change of the envelope surrounding the incoming analog waveform. A nonvarying or slowly varying envelope

indicates steady-state background noise or line noise. A rapidly varying envelope indicates voice. The output of the voice-presence detector is used for the deletion of leading and training silence from messages as will be described later.

Because CCS II customers may signal VSS for control purposes by dialing digits at any time during a call, a DTMF and dial pulse signaling receiver is permanently assigned to each VAC. To overcome the cost penalties of a per-trunk assignment, a new low-cost integrated circuit DTMF and dial pulse receiver has been developed for 1A VSS. The receiver is based on charged-couple device (CCD) filter technology and eliminates the need for bulky RC networks. This new receiver performs at a level of quality equivalent to that of standard central office receivers in approximately an order of magnitude less space and at significantly lower cost. The logic circuitry associated with this new receiver also receives as an input the incoming E-lead signal over which dial pulses from the subtending ESS are repeated. The receiver, upon detecting E-lead activity begins to count dial pulses. Once a valid dial pulse digit is received, the receiver reports the digit as though it were a DTMF digit. Thus, the VAC, PC, and central processor are insulated from the distinction between customer DTMF signaling or dial pulse control. The receiver queues two received digits until the PC interrogates it in order to lessen the scanning load on the PC. The receiver also provides an "early detect" output signal when it begins to observe a DTMF or dial pulse digit as an aid to the record/playback control over messages, as will be described later in greater detail. Once a digit is fully verified, a second "digit present" output is raised to inform the PC that the receiver should be read.

The VAC receives its control via a bit-synchronous serial data transceiver. Each transmission by the PC consists of a 27-bit word, containing 11 address and parity bits and 16 command and status bits. Each transmission to the VAC is followed by a return transmission to the peripheral controller. The bits associated with commands to the VAC are retransmitted back to the peripheral controller in the follow-on reply to allow for transmission error checking by the peripheral controller.

Voice signals to be stored are encoded digitally into a bit serial format at 32 kb/second by an Adaptive Delta Modulation (ADM) CODEC. This CODEC, originally designed for the Subscriber Loop Carrier–40 system, is driven at its level of optimal performance by the AGC circuit which precedes it. Extensive subjective testing has shown that this ADM CODEC with AGC performs at a level of quality comparable with that specified as the Bell System standard for transmission. Upon recording, digitally encoded voice is loaded into a two-port FIFO buffer at the CODEC bit rate. When a sufficient number of bits have been

stored to fill the message storage portion of a disk track (155,078 bits) the disk's controller establishes a path through the network between the buffer and a disk transport and transfers the track's worth of bits at the rate of approximately 10 Mb/second. The buffer must be serviced only once every 4.85 seconds, and hence the disk's controller must be connected to the buffer relatively infrequently. This frees the disk's controller to simultaneously serve the needs of many buffers.

The buffer circuit organizes incoming digitally encoded voice into 1024-bit blocks. Each block consists of 1007 bits of digitized voice, 16 bits generated by the buffer's cyclic redundancy (CRC) generator for the purpose of error detection over the block, and one "voice presence in this block" bit as received from the voice-presence detector. The CRC character is checked by the disk's controller upon transmission to the disk as an error check over the switching path. Upon playback, the disk's controller again checks the CRC characters to verify the integrity of the disk medium. As the message is sent from the buffer to the CODEC, the CRC characters are again checked and stripped out of the bitstream. This again checks the integrity of the transmission from disk controller to buffer, plus the buffer's memory.

Once the buffer is told by the system to begin recording a message sent from the CODEC, it will actually enter a "pseudo record" mode until the voice-presence detector indicates "voice present." In this mode, it records bits but discards them after approximately 120 ms. When the voice-presence detector indicates the presence of voice, the buffer ceases to discard bits and records continuously. This feature eliminates "clipping" of the message potentially introduced by the operational delay of the voice-presence detector.

In the event that a customer should seek to control 1A VSS through DTMF or dial pulse signaling while recording or playback is in progress, the DTMF/dial pulse signaling receiver sends a "freeze" command to the buffer. This prevents the customer from losing any of his message while signaling. The receiver's "early detect" output is sent to the buffer and, when active, causes the buffer to suspend its activity. If a "digit present" signal does not follow, the buffer is immediately "thawed." If a digit is detected, the PC and central processor may elect to force the buffer into the "thawed" mode if appropriate (the customer's digit may, however, be a request for some other action on the part of 1A VSS).

Because of return loss in the trunk circuit and because of transmission echos, a recorded DTMF digit may be fed back into the receiver and be perceived as a customer-initiated signal. To prevent this, the "early detect" signal freezes the buffer, while the receiver continues to listen for DTMF signals. If a DTMF signal is detected while the buffer is "frozen," the signal must have originated from the incoming line and

not from the message being played back. If the signal originates from the recorded message, the receiver suspends its operation for several hundred milliseconds to lessen the frequency of interruption of the recorded DTMF signals.

All VAC circuitry, including the buffer, is implemented on two 8- by 13-inch circuit packs. Thus, the plug-in VAC circuit pack represents the growth circuitry on an engineering basis for expanding offices.

2.3 *The storage media controller*

The functional responsibilities of the disk's controller clearly go beyond that of a computer disk controller. This device, named the Storage Media Controller (SMC), exerts control over the disk media attached thereto, over buffer circuits, and over the electronic switching arrangement that interconnects buffers and SMCs. The SMC is a microprocessor-based controller endowed with considerable intelligence and autonomy. The SMCs communication with the central processor is at a high functional level designed to minimize the frequency of communication with central control and to minimize the real time invested by central control over the detailed operation of its periphery. In this realm, the SMC might receive a work order from the central processor requesting that it play the message beginning at physical location x of its disk media to customer y, who is connected to VAC z. The SMC then accesses the data associated with that message, performs security checks to assure that proper data are being accessed, verifies the integrity of the signal path to the required VAC, and begins to transfer digitally encoded voice to that VAC. This digitally encoded voice is physically realized as many segments of data (see Fig. 3), each segment consisting of 4.85 seconds of recorded voice and 256 bytes of control data. These data include the identity of the customer to whom the message belongs, forward and backward linked list pointers to the next and previous segment of the message, and miscellaneous data regarding the status of the segment. The SMC uses the linked list pointers to manage the playback of an entire message autonomously, and reports back to the central processor when it detects "end of message."

A similar function is executed upon message recording, whereby the central processor furnishes a set of disk storage locations that is available for recording purposes. The SMC manages the building of the data portion of each segment as it is recorded.

The SMC performs somewhat differently in order to compose prompts or system announcements. Digitally encoded voice fragments, each consisting of 63 ms of sound, are "stitched" together by the SMC to create complete phrases. The central processor sends the SMC an ordered set of pointers to many fragments that, taken together, forms a particular phrase. The SMC then seizes the appropriate VAC circuit,

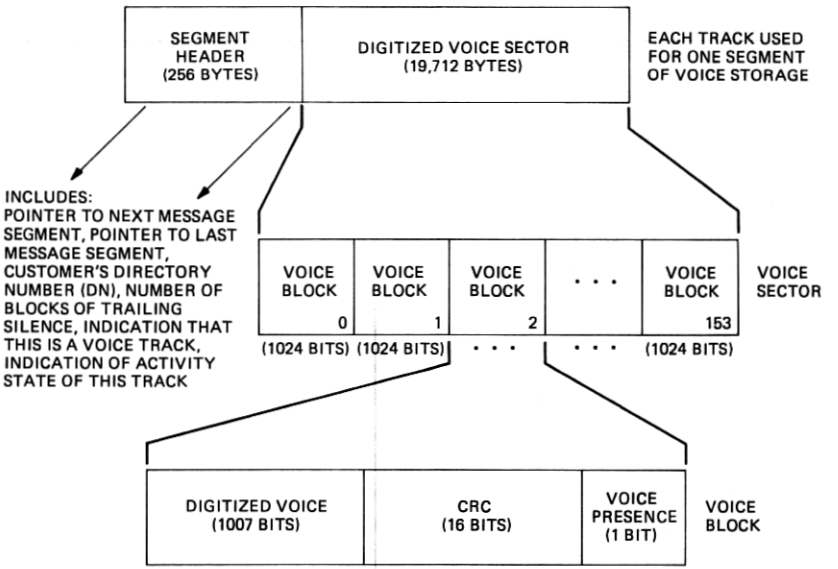

VOICE IS PARTITIONED INTO SEGMENTS, EACH CONTAINING 155,078 BITS (4.85 SECONDS) OF THE TOTAL MESSAGE. EACH SEGMENT IS STORED SEPARATELY ON THE DISK MEDIUM ALONG WITH A DATA HEADER CONSISTING OF A DESCRIPTION OF THE MESSAGE'S SEGMENT POSITION, THE MESSAGE OWNER, THE TYPE OF MESSAGE, AND THE STATUS OF THE MESSAGE.

Fig. 3—No. 1A VSS disk data format.

accesses the fragments from among its disk population, and stores the fragments contiguously in the VACs buffer circuit. The short duration of the speech fragment, taken together with the vast number of fragments that can be stored in the SMCs disk population enable an automatic message composition capability of exceptional quality, variability, and scope. For example, all prompts could be duplicated in several languages, with the target language selected independently for the requirements of a particular call.

The SMC executes trailing silence deletion from recorded messages as a post-processing operation. After a message has been completely recorded, the SMC is instructed by the central processor to autonomously begin with the last segment of a message, inspect it via the stored "voice-presence" bits for the presence of voice, and to delete segments that contain no voice. When the last segment containing voice is identified, its data field is updated to reflect this fact, and to indicate the exact location within the segment of the "end of voice." Upon message playback, the SMC will transmit all bits in segments until the last segment is reached. At that time, only those bits preceding end of voice are transferred to the VAC buffer.

The SMC has the capacity to accept work lists from the central processor to conduct up to 80 concurrent record, playback, or compose

operations. To implement this, the SMC scans the number of VACs assigned to it at any point in time and schedules disk data transfers for those buffers currently in need of service. The VACs that are assigned to a particular SMC at any point in time is a function of system load only. These VACs may be any subset of the total VAC population, and the members of this subset change dynamically with time. For example, if a customer is hearing the playback of six messages in sequence, each of these six messages may have been sent to the VAC associated with that customer by six different SMCs. Thus, that VAC moves from the work list of one SMC to the next as playback proceeds.

The order in which the SMC serves VACs is determined by the locations on disk of the stored segments of data associated with those VACs. The SMC schedules data transfers according to an algorithm which minimizes disk head travel and, thus, maximizes the number of transfers that can be conducted in a given period of time.

In addition to managing the transfer of digitally encoded voice to and from VACs, the SMC serves as a traditional disk controller to serve the central processor's bulk data needs. Data are stored on the same disk media used for voice messages, with special track header fields used to distinguish between locations used for data and locations used for voice. The amount of storage space allocated for voice and for data changes dynamically as a function of system operation and load.

Messages are duplicated to ensure message reliability in the event of a disk transport failure. A message is duplicated after it has been completely recorded and all silence-deletion post processing is complete. A message is always duplicated onto a disk transport associated with an SMC different from that of the original recording. Message duplication is initiated by the central processor causing the initial SMC to seize an idle VAC and to move the first segment of the message to that VAC's buffer. A second SMC is then directed to read the VACs buffer as though it were an incoming message. This process is repeated until the complete message is duplicated. Because the buffer is filled and emptied at the high bit rate associated with buffer to/from SMC transmission, the duplication process transpires at a fraction of the time that was required for the initial message recording.

Each SMC equipped in the 1A VSS may control up to eight disk transports. The fully equipped 1A VSS has eight SMCs associated with it. The disk transports will be described in detail later in this article.

2.4 Switching within 1A VSS

There are two electronic space division switching entities associated with the 1A VSS as shown in Fig. 4. The first is a nonblocking Time-Multiplexed Space Division Switch (TMSDS) which serves to interconnect VACs with SMCs. The second is an engineered blocking electronic

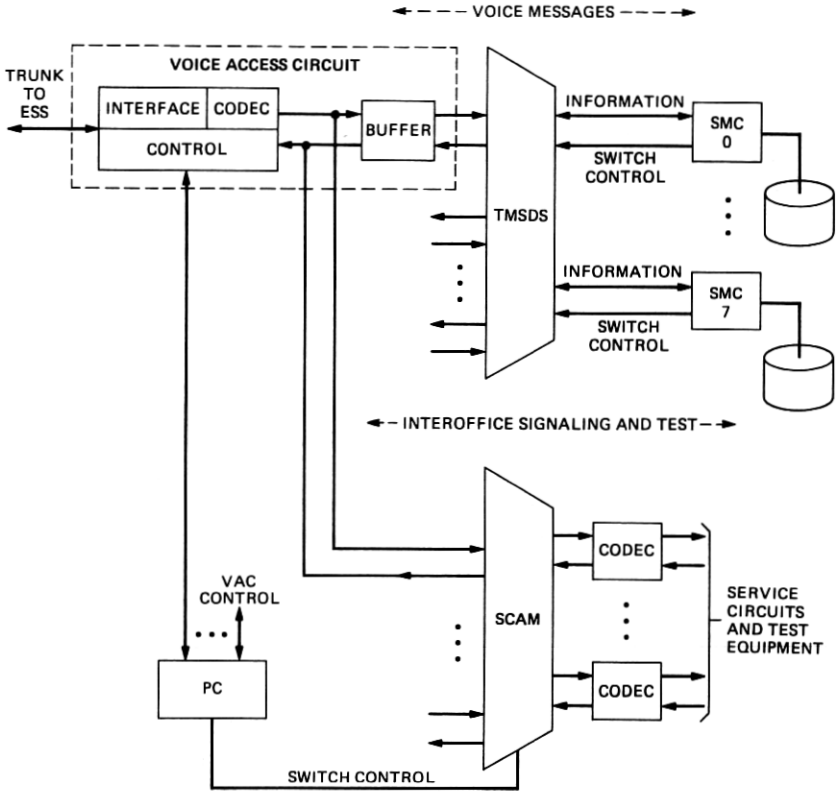

Fig. 4—No. 1A VSS network implemented.

space division switch, called the Service Circuit Access Matrix (SCAM), which serves to interconnect VACs with service circuits and test circuits.

The TMSDS is under the combined control of the ensemble of SMCs. Each SMC utilizes the TMSDS as a time-multiplexed switch with regard to data transfer to/from up to 80 VACs on its work list. Each SMC continuously and sequentially steps through paths to those VACs to interrogate the readiness of each VAC's buffer to conduct a transfer of digitally encoded voice. When a buffer is in a state of readiness, the SMC schedules the transfer according to the current location of the disc track to be used, relative to the current position of the disk's moving head carriage. When the disk is ready for data transfer, the SMC uses the TMSDS to establish a path to the buffer and transfers 157,696 bits (including encoded voice, check bits, and voice-presence bits) in approximately 17 ms. Completing this transfer, the SMC progresses to its next active VAC to effect a similar transfer.

The TMSDS is utilized simultaneously in a manner as described above by as many as eight SMCs. The SMCs operate completely asynchro-

nously with respect to each other and, thus, the TMSDS is under the control of as many as eight independent entities. To prevent collisions of control and/or "lock-up" situations, the TMSDS has integral control circuitry which adjudicates conflicts of control among the SMCs, such as may occur during circuit failures, or during the duplication of voice messages.

A request for access to a VAC by one SMC, while a second SMC is using the VAC causes the TMSDS to put the second request into a "wait" state. The second SMC may wait until an internal time-out occurs at which time it either retries the attempt or determines that a fault has occurred in the TMSDS.

The SCAM is a two-stage space division electronic switch engineered for 0.001 probability of blocking under peak load. The SCAM is controlled by the microprocessor-based peripheral controller in a manner analogous to network control in an ESS. All paths in the SCAM are automatically verified after set-up but before the signal path is completed by a signal loop-around test conducted between the VAC and a "hair pin" in the selected service circuit port.

2.5 The peripheral controller

A microprocessor-based PC has been included in the 1A VSS architecture for the purpose of relieving the central processor of the responsibility for scanning. This peripheral controller also is responsible for the SCAM as previously noted. As described earlier, the PC communicates directly with VACs via a bit-serial communication protocol utilizing a 27-bit word transmitted to and/or from each VAC or SCAM. This word includes commands to the VAC (or SCAM) and received status from the VAC (or SCAM). Additionally, each command sent by the PC to the VAC is acknowledged by the VAC in a response transmission. This provides a checking function to ensure reliability of transmission.

A loop-around transmission is conducted every 40 ms by the PC with each VAC. This serves a basic scanning function associated with trunk status and DTMF signals or dial pulse digit reception. (The DTMF receiver of the VACs queues digits and E-lead signals to permit this relatively slow scanning rate.) Additionally, the PC may be directed by the central processor to perform special communication with VACs at any time in order to effect changes in the state of the VAC (e.g., wink to the ESS, disable the AGC circuit, freeze the recording process). Certain "histories" of events are recorded by the PC for periodic transmission to the central processor. This is exemplified by a history of voice-presence activity seen by each VAC to simplify "time-out-on-no-voice" processing within the application software. Associated with this function, the PC samples the voice-presence detector of an active VAC once each 40 ms and builds a file which documents voice-presence

activity over intervals of hundreds of milliseconds. The central processor periodically interrogates the PC to access voice-presence history files and, thus, is not forced to scan the PC at a high rate to implement a time-out-on-no-voice function.

The PC is duplicated for reliability purposes with one unit active and the other unit in a standby state. The PC incorporates a master clock which controls the population of CODECs. The master clocks of each PC are interconnected in a special arrangement which causes high-speed autonomous switch-over from the clock of the active PCs to the clock of the stand-by PCs upon the failure of the active clock of the PCs. This switch-over precedes the follow-on switch of PCs themselves, which is initiated by the central processor. Thus, a master clock failure does not cause even a momentary system outage because of loss of the CODEC clock. A PC failure is detected by self-check circuitry in the PC and reported to the central processor. The central processor, thus, directs the standby PC to become active. The newly active PC then determines the state of the periphery and assumes control.

2.6 Central processor

The central processor used in the 1A VSS is the Auxiliary 3A Processor (AP). Various versions of this processor are also used in No. 2B ESS, No. 3 ESS, Transaction Network System,[1] and No. 5 Crossbar Electronic Translator System. The AP has the general attributes of a minicomputer with the added reliability of duplicated processors, memory, I/O controllers, and communication buses. Self-checking circuit techniques enable the AP to have rapid fault detection and reconfiguration.

Three types of frame, the AP frame, the supplementary main store frame, and the maintenance frame comprise the central processor for 1A VSS. The AP frame consists of duplicated 3A central controls, duplicated semiconductor memories and I/O controllers. The supplementary main store frame consists of additional semiconductor memory. The 1A VSS uses 768K 18-bit words of memory with the capacity to grow to 1024K words of memory. The maintenance frame consists of a system status panel, two cartridge tape units used for loading programs, and teletypewriter equipment.

2.7 Moving-head disk transports

The storage medium employed by the 1A VSS consists of an ensemble of moving-head disk transports, each with a storage capacity of 300 million bytes of digital memory. This translates into a total voice storage capacity per disk of approximately 21 hours.

The disk media itself consists of 11 platters mounted on a single rotating shaft. Nineteen of the surfaces of these platters are useful for

data storage and each is served by a read/write transducer. The many transducers, or heads, are mounted on arms which extend into the assembly of platters, and the arms are fastened to a common moving-carriage assembly which positions the heads radially on the surface of the platters. The access time to a particular datum is a function of the current position of the carriage with respect to the physical location of the datum on the platter. It is to minimize track access response time that a minimal head travel algorithm was selected to dominate control over job sequencing within the SMC.

The reliability of the 1A vss is influenced by the reliability of the alternating current power which serves the disk transport community. To provide virtually disturbance-free power in the face of the uncertainties of commercial power service, a new uninterruptible power supply (UPS) has been developed for 1A vss. This system, called the TRIPORT,[2] after its internal structure, will react to power failure within a cycle of the disturbance and smoothly convert its own battery-supplied inverter for backup without creating significant power waveform discontinuities.

III. PHYSICAL DESIGN

3.1 1A VSS physical design

The 1A vss equipment uses a standard set of devices, apparatus, and design tools known as the 1A technology. The 1A technology hardware is used in the No. 4 ESS[3] and the 1A Processor,[4] as well as various other ESSs now being manufactured by Western Electric. The use of this technology allowed 1A vss to take advantage of the present manufacturing capabilities of Western Electric. The AP frame, the supplementary main store frame, and the maintenance frame were under manufacture by Western Electric prior to the development of 1A vss.

3.2 Circuit packs

Circuit packs used in 1A vss are FB-, FC-, and FE-coded packs. The FB- and FC-type packs are approximately 4 by 7 inches, with 40 pinouts and 80 pinouts, respectively. The FE-type packs are approximately 8 by 13 inches, with either 80 or 160 pinouts. These packs use the standard WE 946/947-type of connector.

The majority of the circuit packs designed for 1A vss are six-layer multilayer boards with path widths of 8 mils. Where the wiring density of multilayer boards was not required, double-sided and single-sided boards were designed.

Great attention was paid to maximize the density of circuitry per board. For example, the 1A vss VAC is implemented as two circuit packs—the trunk access circuit (TAC) pack and the buffer pack. The

functions located on the TAC pack include the trunk circuit (Type II E&M, 2- or 4-wire), a DTMF receiver, an AGC circuit, a voice-presence detector circuit, a CODEC, and interface and local control circuitry. To achieve this level of density, a new DTMF receiver based on CCD filter technology has been designed for the 1A VSS. Additionally, most digital logic has been implemented as custom-integrated circuits using the new low-power Schotcky gate array technology. Figure 5 is a photograph of the 1A TAC pack.

The buffer circuit pack contains 327,680 bits of random-access memory (RAM) organized as a two-port FIFO memory. Each port may operate asynchronously with respect to the other at speeds defined by the connecting circuits. The 16K RAM used on the buffer pack is the Western Electric coded 28A device. Again, to achieve high-circuitry densities, a large portion of the control logic on the buffer pack is implemented using the gate array technology.

Before the multilayer art masters or the integrated circuit masks for the gate arrays were produced, machine wire-wrapped models were built to test the logic design. Figure 6 shows the buffer pack both in its wire-wrap version and in its final printed wire-board version using gate arrays. Notice that without the use of the custom-designed gate arrays to replace a large number of individual IC devices, the buffer circuitry would have required two packs and connector positions instead of one pack. Because a large number of TAC packs and buffer packs are

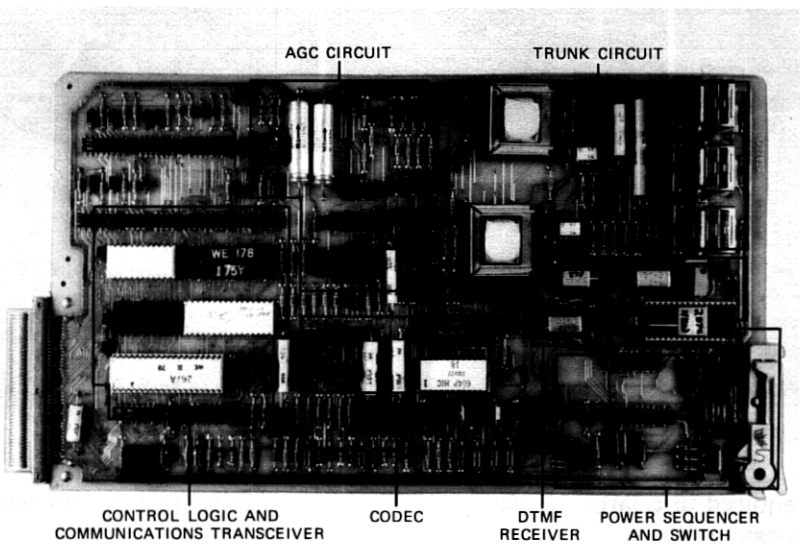

Fig. 5—Trunk access circuit.

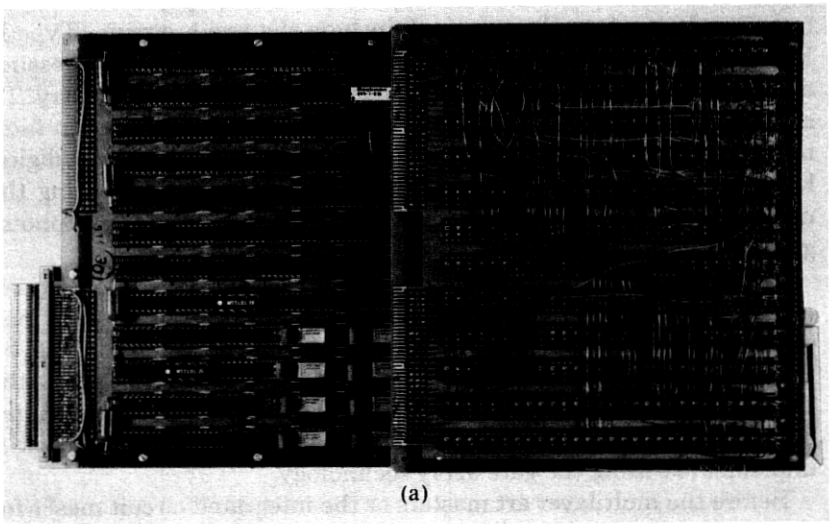

(a)

Fig. 6a—Wirewrap buffer.

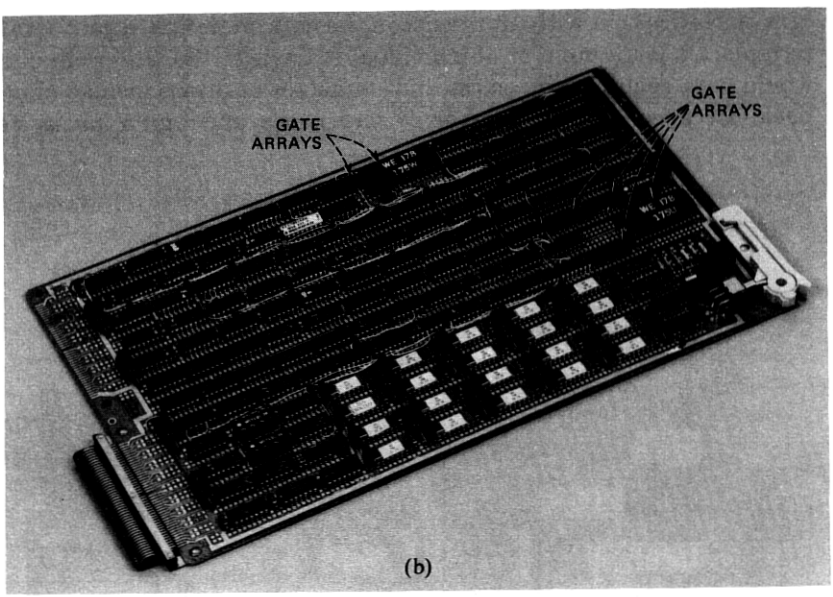

(b)

Fig. 6b—ML PWB buffer.

required in a large 1A vss office, the savings in space and cost were considerable.

All 1A vss circuit packs associated with "growable" functions (for example, the TAC and buffer packs) are designed to be plugged into

and removed from their units without removing unit power. To achieve this, a power switch has been integrated into the circuit pack handle to allow power sequencing of individual packs to be effected automatically. Figure 5, the TAC pack, shows this switch.

3.3 Unit and frame designs

The 1A VSS has been partitioned into functional units, each completely self-contained including power converters and alarm circuitry. Figure 7 is a photograph of the SMC unit used to control the moving-head disk transports. This unit, as well as all of the basic units of the 1A VSS, was designed using the 1A technology. The mounting plate, apparatus mountings, circuit pack connectors, backplane boards and backplane wiring, and designation strips are all standard apparatus used in other Western Electric manufactured units.

The units are mounted on standard 1A-type equipment frames which are 7 ft high by 2 ft 2 in. wide. All frames are 18 in. in depth. Table I lists all of the frames required for a 1A VSS office. The VAC frame is the only frame added to operational systems to serve growth needs. Each VAC frame has a capacity of 32-voice access circuits to the No. 1/1A ESS offices. The minimum number of VAC frames is two to ensure reliability of trunk groups. The VAC frames do not have to be fully equipped. The maximum number of VAC frames is sixteen which provides a total system capacity of 512 TACs. Figure 8 is a photograph of the VAC frame.

Other self-contained equipment provided on an engineered basis are the moving-head disk transports (minimum 3) and the TRIPORT

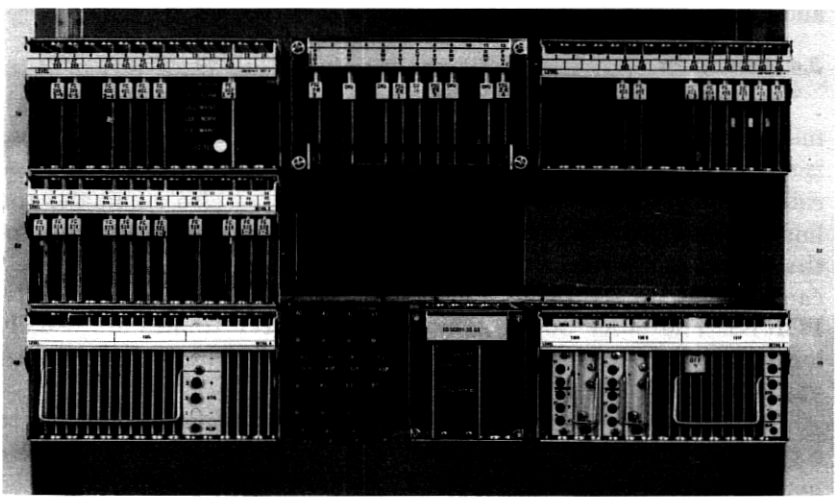

Fig. 7—Storage media controller.

Table I—1A VSS frames

	Size	Number Required Per Office
EXISTING FRAMES		
Auxiliary 3A processor frame	4′4″ (double bay)	1
Supplementary main store frame	2′2″	2
Maintenance Frame	2′2″	1
FRAMES DESIGNED FOR 1A VSS		
Voice access circuit frame	2′2″	2–16
Storage media controller frame	4′4″ (double bay)	1
Peripheral control frame	4′4″ (double bay)	1
Service circuit frame	2′2″	1
Test frame	2′2″	1
Miscellaneous frame	2′2″	1
Power control and distribution frame	2′2″	1
OTHER EQUIPMENT		
TRIPORT cabinet (located in office power plant area)	2′2″	3–8*
Disk transport	23″ x 36″ x 40″ (maximum dimension)	3–24*

* Based on CCS II engineering requirements.

cabinets (minimum 3) which provide the uninterruptible ac power for the disk transports. The TRIPORTS are usually located with the office power plant.

Almost all cabling between 1A VSS units and frames is connectorized. The objective of connectorization is to enable the 1A VSS to be fully assembled, wired, and tested at the factory using the same cables that will be used at site. The only cables not connectorized are the scan and distribute leads associated with alarm circuitry.

3.4 Floor plan

The 1A VSS requires approximately two building bays (approximately 800 square feet) for the maximum-size office. A fixed floor plan is specified in order to use predesigned cabling. This not only ensures cable lengths to meet electrical requirements but also eliminates the line engineering of each cable. Figure 9 shows the floor plan layout of the 1A VSS. As mentioned earlier, the only equipment provided on a capacity basis are the VAC frames, the disk transports, and the TRIPORTS. The office power plant and TRIPORT frames are usually located in a separate power room. Figure 10 is a photograph of a 1A VSS model.

The wiring aisles of the frame line ups are 3 ft wide as compared to a typical 2-ft width in most electronic switching systems. This allows normal office cooling methods to be used, even though a fully equipped VAC frame dissipates approximately 1200 watts.

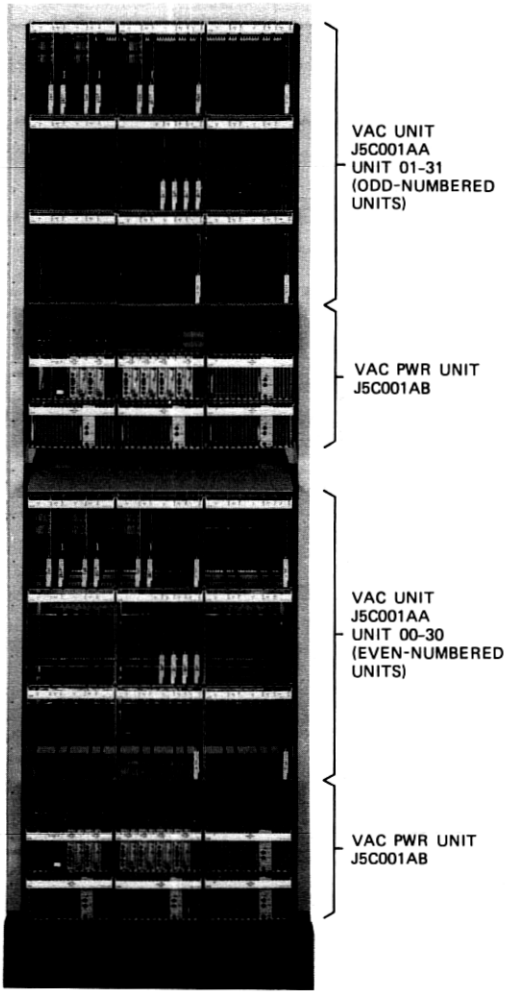

Fig. 8—Voice access circuit frame.

The commercial disk transports used in 1A VSS are designed so that cabling enters the cabinet from the bottom. This is consistent with the normal application of disk transports in computer rooms which use raised floors with all cabling being done under the floor. The 1A VSS has been designed to be located on either a raised floor or on a regular floor. Apparatus is available to support either cabling to the disk transports under a raised floor or in overhead cable raceways.

IV. CONCLUSION

A new centralized system has been developed for the purpose of

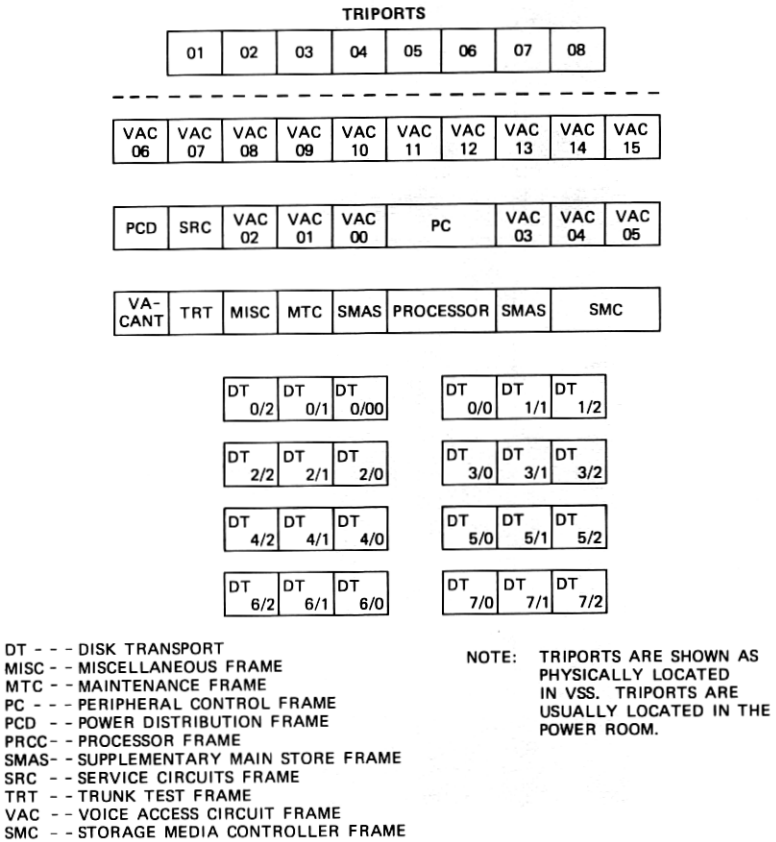

Fig. 9—Voice storage system floor plan.

adding high-capacity voice storage and voice storage processing capabilities to the SPC network. This system has been developed using a mixture of new periphery design and use of an existing processor complex. The system has been implemented in a generalized fashion

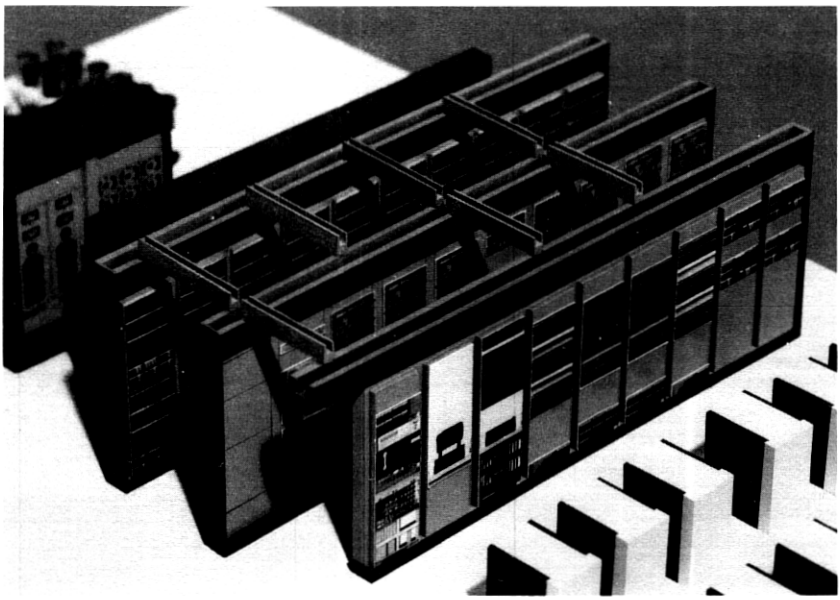

Fig. 10—Model of 1A VSS.

which allows the addition of future storage services with minimal (if any) impact upon the system hardware architecture.

REFERENCES

1. Transaction Network, Telephones, and Terminals, B.S.T.J., *57,* No. 10 (December 1978), pp. 3325–539.
2. R. Ostapiak and R. Rando, "Triport a Low Cost Source of Continuous AC Power," Bell Labs Record (July–August, 1978), p. 175.
3. J. H. Huttenhoff et al., "No. 4 ESS Peripheral System," B.S.T.J., *56,* No. 7 (September 1977), pp. 1029–55.
4. J. O. Becker et al., "1A Processor: Technology and Physical Design," B.S.T.J., *56,* No. 2 (February 1977), pp. 207–36.