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

1A Voice Storage System:

Software

By G. W. GATES, R. F. KRANZMANN, and L. D. WHITEHEAD

(Manuscript received July 31, 1979)

The new Custom Calling Services II (CCSII) have been provided by adding a 1A Voice Storage System (1A VSS) as a new node in the Stored Program Control network. Software and a new trunk circuit are required in the No. 1/1A ESS to provide call control, call filtering, and routing to a 1A VSS. The 1A VSS accepts the call and provides the package of voice services known as CCS II. The software required to provide these services is described.

I. OVERVIEW

1.1 Design considerations

The software required to implement the new line of Custom Calling Services (CCS II) being developed for the Bell System exists in both No. 1/1A ESS and in the 1A Voice Storage System (1A VSS). The software in No. 1/1A ESS is required for call screening, for determining which calls should receive service by 1A VSS and for dealing with the interaction of existing services with this new class of services.

There is a strong interaction between the ESS and 1A VSS in providing CCS II services. To the extent possible, all service control has been placed in the 1A VSS. In addition, all customer data and control data are maintained on the disks in 1A VSS. In a general sense, as soon as ESS determines that the calling party requires a 1A VSS-provided service, the call is routed to 1A VSS. All control data are permanently maintained in 1A VSS and required data are sent to ESS when it is needed. When the data are no longer needed, as when a call answering customer deactivates,[1] ESS destroys the data so as to regain the

memory space; the data are retransmitted when the customer next activates.[1]

1.2 The SPC network

The partition of function between the ESS and the 1A VSS subsystem is an example of the growing trend in the Bell System to specialize functions in the various nodes of the Stored Program Control (SPC) network. The SPC network is the name applied to the collection of stored program controlled systems which provide customer services and Bell operating companies (BOCs) administrative services. These systems are interconnected by trunks and data links and, hence, are referred to as a network. The network includes the increasing number of Electronic Switching Systems (ESSs), the Operations Support Systems (OSSs), the Traffic Service Position Systems (TSPSs), and 1A VSSs. This growing network of stored program controlled systems permits increased sophistication in customer features and in techniques for providing these features. The 1A VSS services are an example of the concentration of customer feature implementation in a specialized type of node in the SPC network. Such a node can provide its specialized features to many class 5 offices by having calls routed to it for service.

Figure 1 illustrates the role of 1A VSS as a node in the SPC network. The connection to class 5 No. 1/1A ESSs and to OSSs is shown.

II. SOFTWARE IN NO. 1 ESS

2.1 Partition of functions between the host office and 1A VSS

An early objective in the design of 1A VSS features was to minimize the impact on the interconnecting (host) ESS and to place the major burden of responsibility on the 1A VSS itself. There were several reasons for this important principle:

(*i*) The 1A VSS processor and system software were new and, hence, provided flexible vehicles which could more readily support functional changes as the system matured.

(*ii*) Although the 1A VSS was initially to serve the No. 1/1A ESS host office, other host systems such as No. 2B ESS, No. 3 ESS, No. 5 ESS, and No. 5 Cross-bar Electronic Translater System were also considered as potential candidates for the future. Any function performed within the 1A VSS would need to be developed only once, whereas each host office function would require separate development on each system.

Figure 2 depicts the major functions performed by the host ESS and the 1A VSS. The 1A VSS itself provides the capability for recording, storing, and returning voice messages and announcements, and for interacting directly with the customer to provide the services described in the companion article.[1] For the customer to make use of these

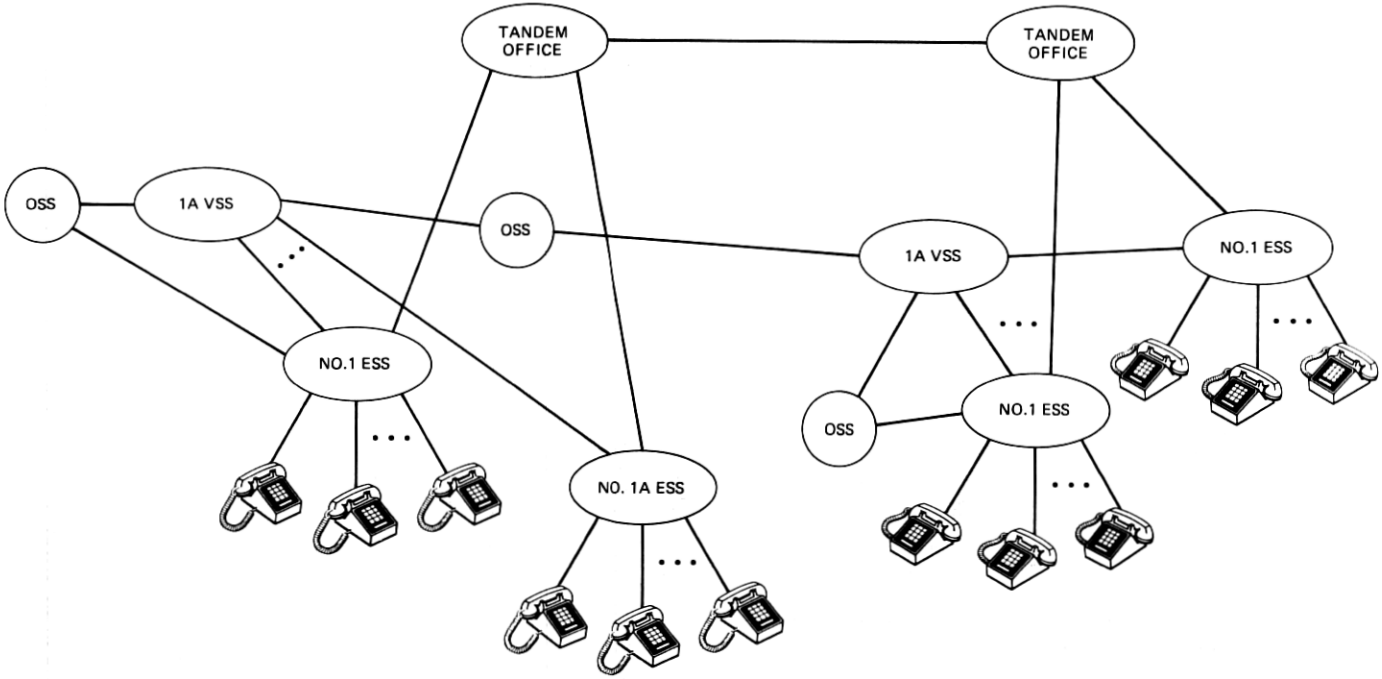

Fig. 1—1A VSS: a new node in the SPC network.

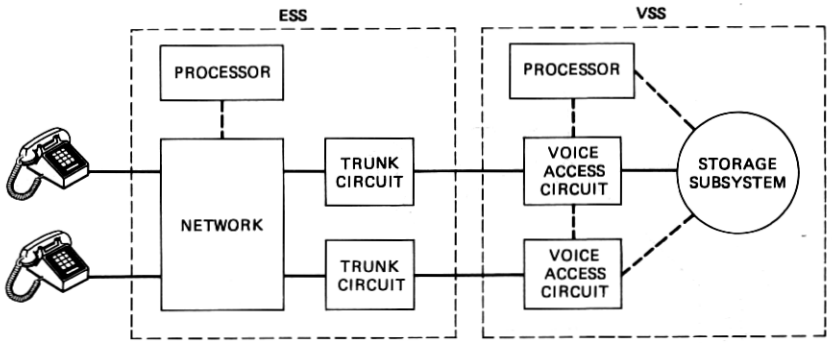

ESS

- DETECT REQUESTS FOR VSS SERVICE
- PROVIDE CUSTOMER ACCESS TO VSS
- MAINTAIN STATUS OF ACTIVE CUSTOMERS
- INTERCEPT CALLS AND ROUTE TO VSS
- PROVIDE MESSAGE WAITING ALERTING
- PROCESS SERVICE ORDERS AND PASS DATA TO VSS
- PROVIDE LOCAL VSS SERVICES
- PROVIDE SMOOTH INTERACTION WITH EXISTING SERVICES
- MAINTAIN ESS/VSS TRUNKS

VSS

- ANALYZE AND PROCESS SERVICE REQUESTS
- STORE, RETRIEVE, AND PLAYBACK MESSAGE SEGMENTS
- PROVIDE FLEXIBLE ANNOUNCEMENT CAPABILITY
- MANAGE DISK SPACE
- PROVIDE FOR INTERACTIVE CONTROL BY CUSTOMER
- ACCUMULATE DATA FOR BILLING AND TRAFFIC
- INITIATE CALLS THROUGH ESS
- MAINTAIN VSS HARDWARE

Fig. 2—ESS/1A VSS software functions.

services, the host ESS provides several capabilities, some new and some extensions of existing capabilities. Essentially, these capabilities are of two types: those that are general and those that are related to specific features. General capabilities are as follows:

(*i*) Customers gain access to 1A VSS by dialing the special service prefix (* or 11) plus two digits, or by dialing seven digits.

(*ii*) ESS and 1A VSS processors communicate using an expanded form of multifrequency signaling.

(*iii*) Service orders are entered from the ESS with subsequent transmission of service order data to 1A VSS.

(*iv*) Audits of new and modified data are performed to assure the integrity of transient data.

(*v*) Maintenance of the ESS trunk circuit and the transmission facility to the 1A VSS is provided.

(*vi*) Resource usage counts are recorded for the software resources used.

Capabilities that are related to specific features include:

(*i*) Terminating calls are intercepted and rerouted to the 1A VSS. Intercept is of three types: immediate, busy, and don't-answer. Call Answering service typifies the use of the intercept capability within the ESS.

(*ii*) The customer is provided an indication that the 1A VSS has voice messages waiting to be retrieved by the customer.

(*iii*) Voice messages are delivered from the 1A VSS office through the host ESS for the Advance Calling feature. Also included is the handling of a special AMA billing message sent from 1A VSS for billing the terminating portion of the Advance Call.

(*iv*) Capability is provided for a customer to monitor a call being recorded and to answer the call personally if desired. This feature is called Monitor/Cut-Through.

(*v*) Coordinated interaction of the new CCS II features with existing features, including those of CCS I, is provided.

The following section presents an overview of the software required for the host ESS in order to implement these capabilities.

2.2 ESS software design overview

The host ESS for the initial implementation of CCS II is the 1/1A ESS. It is beyond the scope of this paper to discuss the detailed structure of the 1/1A ESS implementation for 1A VSS services since it would require substantial background in the design of 1/1A ESS software and hardware. Hence, the design will be presented conceptually, allowing the reader to mentally apply it to any familiar switching system, as appropriate.

The ESS software for implementing the new Custom Calling Services can be viewed as a set of new and modified capabilities, each providing a particular part of the service. Figure 3 illustrates those capabilities and the control flow among them. The following conventions apply in Fig. 3:

- Circles represent input/output devices as follows:
  - SO TTY—Service Order (SO) Teletypewriter.
  - LINES—Both 1A VSS customer lines and others.
  - TRUNKS—Both interoffice trunks to the 1A VSS and to other switching offices.
- Rectangles represent the various functional capabilities.

The primary function of each 1A VSS-related capability is discussed below.

The Service Order Handler accepts messages from the SO TTY, screens messages to ensure that they are syntactically correct and that the subscriber specified by the message is permitted access to the services specified. It also assembles appropriate SO messages to send to the 1A VSS via the Data Message Sender. All 1A VSS subscriber service orders pass through a host ESS prior to transmittal of the service order data to the 1A VSS processor. This is done primarily to consolidate the administrative aspects of service order processing and to allow the ESS to make appropriate screening checks. Thus, a service

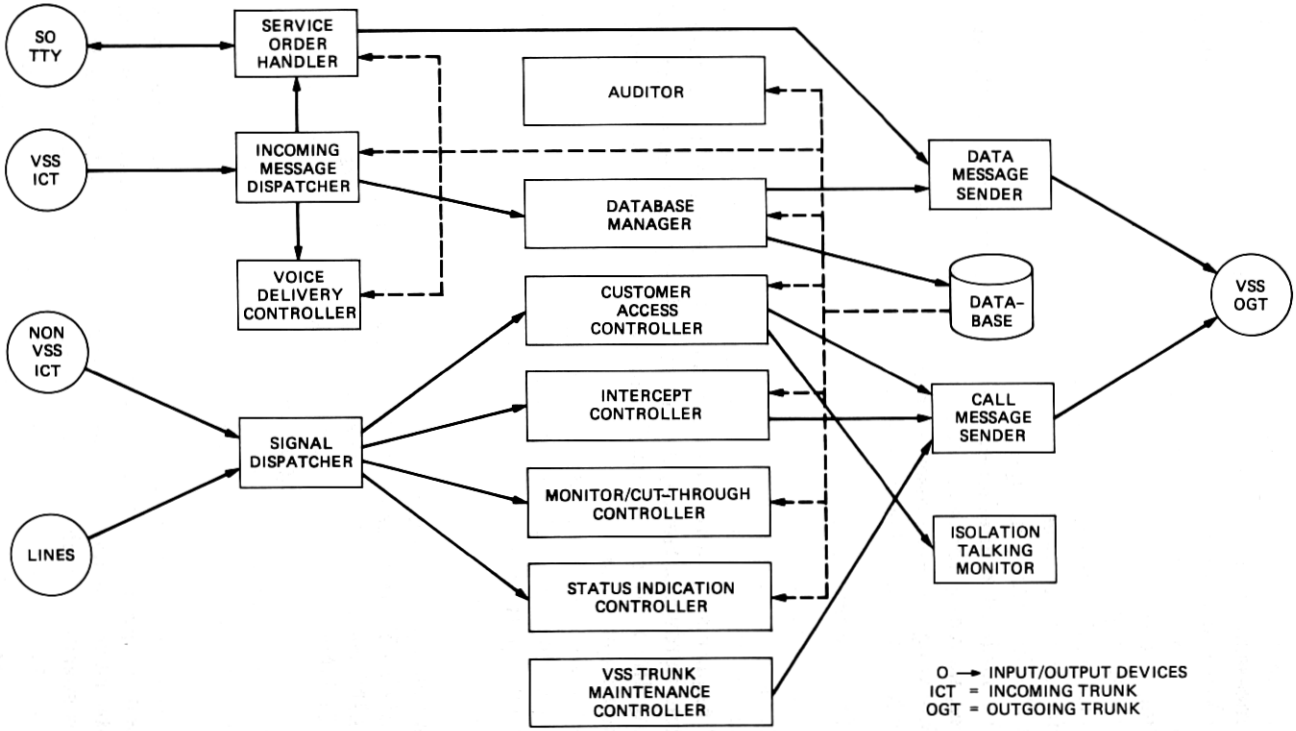

Fig. 3—Host ESS conceptual design.

order describing a new customer for the ESS can simply identify the 1A VSS services that the customer wishes to purchase as part of a single order. Note that the entire customer profile describing all aspects of the 1A VSS service purchased is maintained in the 1A VSS processor and that those items required by the ESS processor are sent to the ESS from 1A VSS when the subscriber activates the service.

The Service Order Handler also is responsible for processing 1A VSS customer-related VERIFY messages which allow the ESS craft person to check the content of the customer profile in the 1A VSS processor. This requires the transmittal of a VERIFY request from ESS to 1A VSS and a VERIFY RESPONSE message in the opposite direction.

The Incoming Message Dispatcher receives all MF message packets from 1A VSS trunks, analyzes the dispatch code, and distributes the data to the appropriate client program.

The Voice Delivery Controller provides functions required for the delivery of Advance Calling messages. This essentially resolves into a terminating call (if the destination telephone is in the host ESS) or into a tandem call (when the destination is not in the host ESS), with the 1A VSS incoming trunk serving as the originating terminal in both cases.

The Signal Dispatcher is a collection of routines which interprets signals from lines and trunks to decide which other processes should be requested to further process the signals. For example, when a customer dials the Call Answering access code *51, the Signal Dispatcher interprets the *51 as a request to activate Call Answering service and passes control to the Customer Access Controller.

The Auditor represents the collection of audit programs which assesses the consistency of data for 1A VSS-related features. Audits in ESS systems form a powerful force to maintain the stability of the host ESS. Several new data structures for 1A VSS services have required corresponding new audit software, while extensions of existing data structures required modification of extant audits.

Note that only the Database Manager is allowed write access to the database representing the pertinent customer profile data within the ESS for active subscribers, whereas many other capabilities are given read-only access. The database for an active Call Answering subscriber includes such items as message-waiting tone and message-waiting ring indicators, Monitor/Cut-Through feature allowed, and identity of the trunk group to the 1A VSS. Customer service requests that require access to the 1A VSS are handled by the Customer Access Controller. It screens the request to assure that the request is allowed within the ESS, formats an appropriate MF packet, selects a trunk to the 1A VSS and passes control to the Call Message Sender. Since some CCS II services are offered on a usage-sensitive basis within the host ESS,

potentially all lines within the office can request activation of these services. However, some combinations of usage-sensitive services, such as Call Answering on COIN lines, are inconsistent or confusing. Therefore screening, based on the originating and terminating major class of the line, is used to control such situations so that specific services may be denied to particular line classes.

Whenever a CCS II subscriber has an active intercept feature, such as Call Answering the Intercept Controller assumes control of any call that would normally terminate to the subscriber's lines. Its function is to perform the actions necessary to route the call to the 1A VSS.

The Monitor/Cut-Through Controller processes requests for this subfeature of Call Answering service and is illustrated in the sequence shown in Fig. 4. The MONITOR function is accessed when the Call Answering subscriber dials the appropriate two-digit access code. A check of the subscriber's database is first made to determine whether the Monitor/Cut-Through feature is allowed.

The Status Indication Controller provides for Message Waiting Tone (MWT) and Message Waiting Ring (MWR) to alert a Call Answering subscriber that messages are waiting to be retrieved from the 1A VSS. Both status indications are under control of the customer's database profile established by the service prototype (see section on the 1A VSS Feature Processor Subsystem). The MWT is provided when applicable on all call originations using software control of dial tone through standard digit receivers. The MWR is a short burst of ringing applied to the customer's line following disconnect from stable network connections involving that line.

The Call Message Sender and Data Message Sender perform the task of transmitting MF packets of information to 1A VSS for various software clients as described in the previous section. These modules perform functions, such as seizing appropriate memory resources, a 1A VSS trunk, and an MF transmitter, as well as implementing the interprocessor communication protocols.

The Isolation Talking Monitor disables the Call Waiting feature and any similar features that may apply tones or switching noise to a customer's line. This capability is invoked principally, while the customer is in the process of recording a greeting or message on the 1A VSS, but it may be utilized in other circumstances to prevent adverse interaction of CCS II with other services to which the customer may have access.

The 1A VSS Trunk Maintenance Controller is responsible for providing new diagnostic software to verify the operation of the new two-port trunk circuit, as well as the standard end-to-end operational test provided on many interoffice trunks. Additionally, provision to allow the standard transmission quality tests both automatically and man-

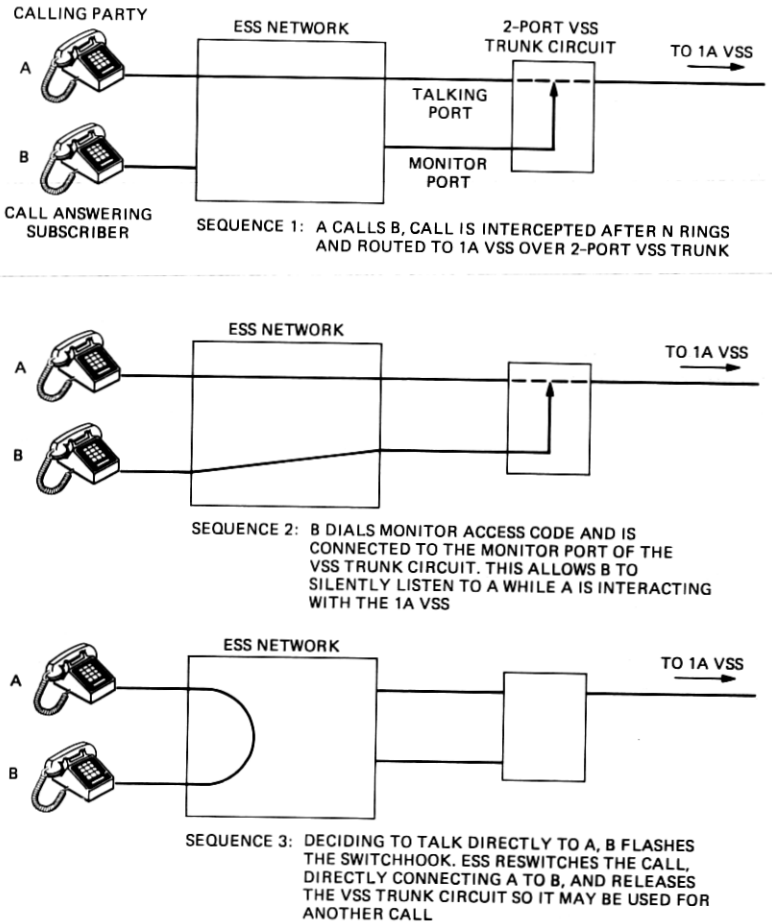

Fig. 4—Monitor/Cut-Through capability.

ually from the various test panel configurations is contained within the Trunk Maintenance Controller conceptual model.

2.3 Interaction with existing customer services

Great care was taken to assure that the new CCS II services mesh well with the many customer services already provided by the 1/1A ESS. Two examples are given here to illustrate this interaction.

Example 1—Call Answering interaction with Call Forwarding.

Both of these services allow the customer to accomplish a similar goal, namely, when activated, each will result in calls that would normally terminate at the subscriber's line being routed to an alternate

destination. For Call Forwarding, the customer specifies the destination; for Call Answering, the new destination is the 1A VSS. Obviously, both cannot be active at the same time or the intent would be ambiguous. Hence, only one type of intercept service is allowed to be active at any one time. If one service is active, subsequent attempts to activate the other result in reorder tone being applied to the customer line.

Example 2—Call Answering (busy line) interaction with Call Waiting.

In this case, a priority of action is used. Call Waiting is useful in informing a customer whose line is busy that another caller is trying to reach the customer. This is done by applying a short beep-tone to the customer's line at intervals of 10 seconds. The customer, wishing to answer the new incoming call, verbally informs the original party of his intent to do so. He then flashes the switchhook which results in the customer being connected to the new party, while the original party is placed on "hold." A subsequent switchhook flash will bring back the original connection. However, if the customer also has Call Answering service active and elects not to answer the new caller, the latter will be intercepted after the first 10-second period and will be routed to the 1A VSS. In this way, the Call Waiting and Call Answering services conflict minimally and provide enhanced call control capability for the 1/1A ESS customer.

Feature interactions of this sort are implemented wholly within the host ESS.

2.4 The Interface between VSS and ESS

As shown previously, 1/1A ESS customers gain access to the 1A VSS via two-way voice frequency trunks interconnecting the two systems. Signaling associated with the use of these trunks is accomplished via an expanded form of E and M, Multifrequency (MF), wink-start signaling, which is still the most common interoffice signaling arrangement in the Bell System. This communication arrangement was selected since it was most easily adaptable to other existing switching systems.

Communication messages can be divided into two main categories: (*i*) those messages which will normally proceed to a talking connection between a customer and the 1A VSS—Call Messages, and (*ii*) messages which involve transmission of call control data only—Data Messages. Signaling protocol for Call Messages is quite standard, except for the content and amount of information to be transmitted. Typical information in an MF Call Message packet specifies the type of call message (e.g., Call Answering activation, or Call Answering intercept), the restriction class, the 1A VSS customer identity (by Directory Number),

and the billing specifications. Standard Call Message protocol is illustrated in Figure 5, Section A.

Consider for example the Call Message MF packet requesting activation of Call Answering service for a casual user. The MF packet sent from the host ESS to 1A VSS would be of the form

	MESSAGE DISPATCH CODE	CUSTOMER DN	RESTRICTION CLASS	BILLING DN	
⟨KP⟩					⟨ST⟩,

where

KP =KEYPULSE DIGIT

Message Dispatch Code =2-digit code identifying the receiving client program in 1A VSS.

Restriction Class =For example, "complaint observing requested by this customer."

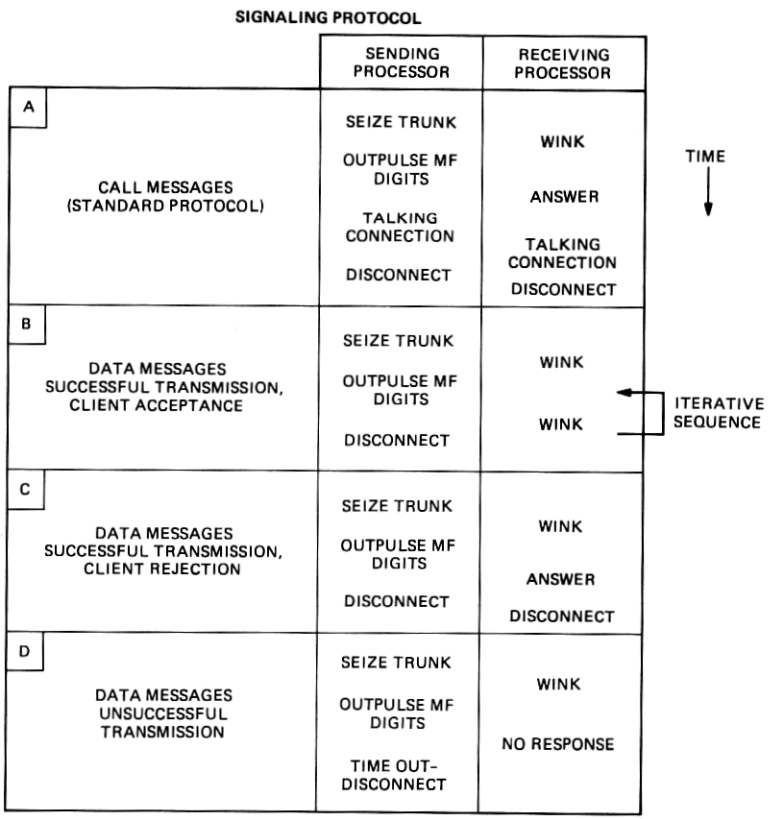

Fig. 5—ESS/1A VSS interprocessor communication.

Customer DN	=4-, 5-, or 7-digit form of the customer directory number (DN).
Billing DN	=Directory number to which 1A VSS should bill the charges for this use of Call Answering service (optional).
ST	="START PROCESSING" digit (end-of-message).

There are 99 two-digit message dispatch codes. Data Message MF packets have much the same format as Call Message packets. The dispatch code distinguishes the packet as a Data Message and identifies the client program in the receiving processor. Examples of Data Messages are (*i*) the activation message sent from 1A VSS to ESS in response to an ESS customer's activation request, and (*ii*) service order messages and their replies.

Signaling protocol for Data Messages is necessarily more complex than the Call Message protocol since it involves interprocessor communication without the presence of the customer to detect the success or failure of the communication. Figure 5, Sections B, C, and D depict the signaling protocol for Data Messages. Note that one of three responses is expected from the receiving processor:

WINK—Implies successful transmission of the message and acceptance by the client program.

ANSWER—means the MF packet was received but was rejected by the client program, for example, because of incorrect format.

TIME-OUT—implies unsuccessful transmission in the same sense as standard interoffice MF signaling.

To increase trunk usage efficiency for data transmission, capability to batch Data Messages is provided. Batching means that multiple independent Data Messages may be transmitted over a single trunk. After receiving the WINK acknowledgment, the transmitting processor will either disconnect, signifying end of transmission, or commence sending the next Data Message.

III. SOFTWARE IN THE VOICE STORAGE SYSTEM

3.1 Software techniques

Control in 1A VSS is distributed among the central processor and several microprocessors which control the periphery. The microprocessors provide the necessary low-level repetitive control and relieve the central processor of this workload.

A major goal in the design and implementation of the central processor software was that it be easily modifiable. Several techniques and rules were used to achieve software that allows new features to be added easily.

(*i*) The fundamental technique is embodied in the software architecture. The software architecture was designed to partition the complexity of the system so that designers and programmers have to concern themselves with only a subset of the total problem.

(a) Software that requires knowledge about the detailed hardware characteristics is concentrated into a few subsystems and the need for this knowledge is eliminated from other subsystems. This technique makes feature development easier by controlling the need for detailed hardware knowledge and makes it easier to change the hardware and firmware. It also results in less overall software impact when hardware and firmware are changed.

(b) The software that controls feature operation is concentrated in one subsystem. Other subsystems provide high-level service functions to the feature subsystem. This effectively creates a language of functions which can be used to build and expand services.

(*ii*) Within the feature subsystem each feature is implemented as independent software. To do otherwise would mean that each time one feature was changed, other features could be affected.

(*iii*) Separate data structures are built for each feature. Shared data structures for one customer seems natural, but if the features are completely disjoint, then the data are kept disjoint to avoid interaction effects.

(*iv*) A high-level language with data structure definition capability is used.

(*v*) Many characteristics of 1A VSS operation are implemented as system parameters so that they will be easy to change as experience is gained from early customer use. Some of these parameters will require software recompilation to change, while others can be changed by modification to the system in the field. Examples of system parameters include: (a) number of seconds of silence before time out during recording, and (b) the time to return answer supervision during a call. The concept of parameters is oriented toward overall system characteristics. An analogous concept of options on specific customer features is used and is discussed in the section on feature implementation.

(*vi*) The software was built with the rule, "Design it correctly, build it, then tune it." Tuning a system too early can destroy its structure and, hence, destroy its modifiability.

(*vii*) The call-related portions of the software were designed and implemented using finite-state-system concepts.

The finite-state-system design technique consists of partitioning the software into functional models where each model is viewed as a finite state automaton. The model consists of states, signals, and transition

routines. The occurrence of an event causes a signal to be sent to a model which is in a particular state. The model executes particular transition routines as a function of its state and the received signal. It then enters another state to wait for another signal. Figure 6 illustrates the diagram of such a model.

Finite-state design techniques provide a good structure for the implementation of call processing. They produce software which is self-documented by the state diagrams of each model. Because of this documentation and the intuitive naturalness of the structure, the resulting software is easy to read and understand.

These techniques have produced call processing software which should be easy to modify as new features are added to the 1A vss. Initial indications are that this goal has been met, but several years of experience will be required before a final judgment can be made.

3.2 Software architecture overview

The 1A vss software runs under control of the Extended Operating System (EOS), a real-time control system developed for Auxiliary 3A Processor applications. Call handling software runs in a single EOS task and is controlled within the task by the State Table Controller (STC). The STC provides the structure necessary to process signals and to

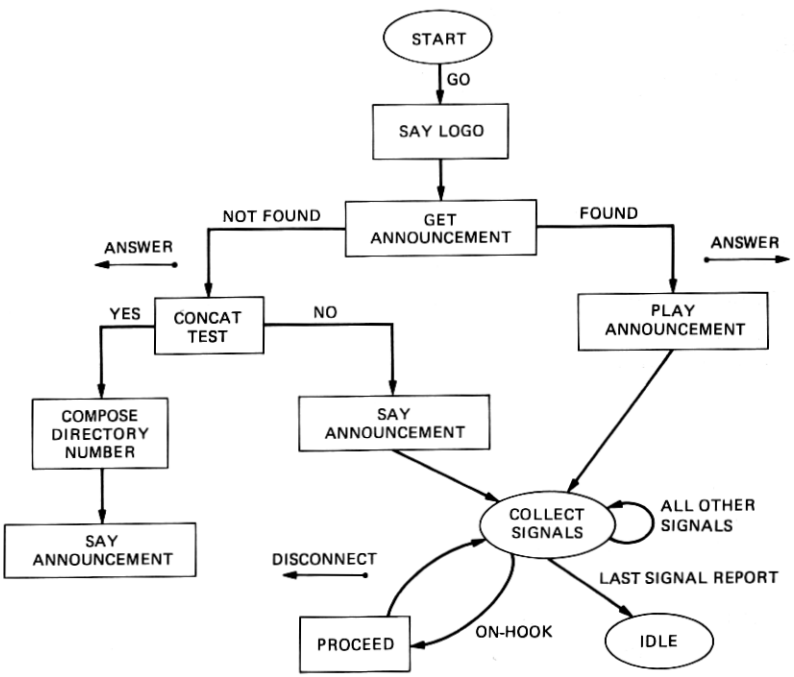

Fig. 6—Example of finite state model.

control models as required for the VSS finite-state design techniques. The STC schedules models, queues signals and maintains state control for each model.

The overall structure of the call processing software is shown in Fig. 7. The software is divided into six major subsystems. The Feature Processor controls the actual customer features. It calls on the Recording Trunk and the Database Manager for services; these systems in turn request lower-level services from the Input/Output Processor, the Voice Manager, and the File System. Support services are provided by device handlers, disk memory allocation software, a message duplication service, and disk erasure software.

3.3 The feature processor subsystem

The characteristics of a customer feature are incorporated in the Feature Processor subsystem. Like all call processing software, the Feature Processor is a collection of finite state models which performs transitions from state to state as the various call events occur. Events such as "off-hook," "customer dialed a 6," or "message playback complete" cause signals to be sent to the appropriate model. The model executes transition subroutines, sends signals if required, and enters another state to await the next signal. Each call in the system creates an "instance" of each model as the call progresses (similar to a software process). Multiple call capability comes implicitly from the collection of all instances of these models.

The set of Feature Processing models and associated transition routines orchestrates the handling of calls but does little of the actual work. The work is done by calling on the Recording Trunk and the

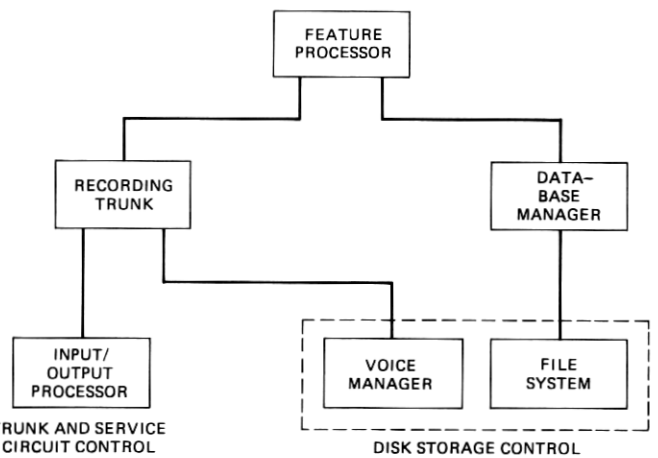

Fig. 7—Basic structure of call processing software.

Database Manager. These two subsystems provide an extensive set of high-level functions which constitutes a primitive language for building customer services.

Examples of the type of functions provided include:

(*i*) Report origination
(*ii*) Report on-hook
(*iii*) Return answer
(*iv*) Get dialed digits
(*v*) Disconnect
(*vi*) Seize trunk
(*vii*) Send data
(*viii*) Say a system announcement
(*ix*) Play a customer message
(*x*) Compose an announcement from fragments
(*xi*) Record a message
(*xii*) Erase a message
(*xiii*) Secure customer data
(*xiv*) Release customer data
(*xv*) Modify customer data.

Many options have been incorporated into each service in order to be responsive to the changing needs of the telephone customer.

The solution to handling the changing needs of the Call Answering customer was to implement essentially all the options which were considered useful and to provide a way to define a customer feature as a collection of these options. Additionally, several packages of Call Answering services can be marketed by defining several collections of options. A set of options is called a prototype, thus, a package of Call Answering options is defined by defining a prototype. The three packages of Call Answering to be marketed initially, i.e., Daily, Monthly, and Deluxe, are created by defining three prototypes with the associated, required sets of options. Further flexibility was gained by providing the capability, through customer service orders, to modify each of the options for the individual customer. Thus, a deluxe customer can have the maximum length of a message extended from 30 to 60 seconds by a service order indicating such a change for that one customer.

IV. THE RECORDING TRUNK SUBSYSTEM

The Recording Trunk Subsystem is an abstraction of an "idealized voice storage trunk." Such an "idealized" trunk in 1A vss would be capable of recording and playing messages and handling timed sequences in an autonomous manner. By abstracting these characteristics and incorporating them into a software system, feature designers

are given a powerful capability for building voice features which do not require intimate knowledge of the complex 1A VSS architecture. The Recording Trunk Subsystem provides the feature programmer the ability to:

(*i*) Record a message
(*ii*) Play a message
(*iii*) Erase a message
(*iv*) Return answer supervision
(*v*) Control silence timeout
(*vi*) Acquire allowable digits
(*vii*) Control digit timing
(*viii*) Recognize flash signaling
(*ix*) Send messages to the ESS office
(*x*) Receive messages from the ESS office.

The Recording Trunk calls upon the Voice Manager (VM) and the Input/Output Processor (IOP) in providing functions to the Feature Processor.

V. THE DATABASE MANAGER AND FILE SYSTEM

Data services are provided to the system by the Database Manager and the File System. The 1A VSS Database Manager was designed and tuned to fit the type of support needed by VSS features. Rapid access to the customer database is provided by the physical clustering of logically adjacent data. Flexible database services are achieved by basing the design on the general concepts of the relational model of data structures.

The File System provides for the random access storage and retrieval of variable length records. To provide the required reliability, each record is duplicated when written. The File System and the Database Manager are designed to specifically complement each other so as to meet the objective of minimization of data storage and transfer costs. The File System also provides storage services directly for administrative data such as billing and traffic data and the collection of data on equipment failures.

VI. INPUT/OUTPUT PROCESSOR

The IOP provides functional control and error detection for the 1A VSS trunks and service circuits. In this capacity, it receives requests for service from the Recording Trunk and system maintenance software. Control is achieved through interaction with the microprocessor-based peripheral controller, with responses from the periphery distributed to the requesting subsystem. High-level functional device requests are

accepted by the IOP and transmitted to the periphery as a sequence of device commands.

The IOP also receives notification of autonomous events from the periphery, e.g., a trunk seizure. These are distributed, as appropriate, to the associated Recording Trunk, maintenance or error control software. Timing control and error recovery are also provided by the IOP.

VII. THE VOICE MANAGEMENT SUBSYSTEM

7.1 Voice manager

The Voice Manager encompasses all software for control and manipulation of stored voice. The operational unit of access is the message. Messages may have a variable length and are comprised of one or more fixed-length segments. These segments are the fundamental units of storage allocation and deallocation. Each message has a unique owner. The owner may be either a customer, designating a message entered in conjunction with the customer's service, or the owner may be the system, designating a system announcement identified with a particular VSS service.

The Voice Manager provides the Feature Processor three basic capabilities for manipulating messages: the ability to record a message, the ability to play a message, and the ability to erase a message when no longer needed, thus, releasing the space for other uses. Each service is defined as a sequence of these operations with appropriate system announcements played to prompt the customer.

Because it is impractical to store certain system announcements in prerecorded form, e.g., "You have *seven* messages," the VM provides the capability to play such messages from a small set of prerecorded fragments. This allows the Feature Processor to specify the phrase "You have seven messages" as follows:

Specification	Fragments
Fragment identifier	"you have"
Decimal number	seven
Fragment identifier	"messages."

To allow a reasonable range of numbers, individual fragments are recorded. The following are examples of such fragments: the numbers 1 through 19, plus 20, 30, such phrases as a.m., p.m., and the names of the days of the week.

7.2 Storage media controller handler

The Storage Media Controller (SMC) handler provides the communication path between the Voice Manager and the microprocessor-based peripheral device which performs the voice commands.

The structure of the handler is governed by two characteristics of the SMC: its ability to service many active calls simultaneously and its highly autonomous operation. The handler defines four phases of operation, described below, and provides all necessary synchronization. These phases are staggered for the equipped SMCs to smooth the main processor load.

Command Phase—The handler transmits all voice commands accumulated during the last cycle.

Voice Phase—The handler is dormant. The SMC schedules and performs the commands it has been given, signaling the handler when complete.

Reply Phase—The handler reads the status replies indicating the disposition of the commands performed during the voice phase.

Disk Controller Phase—The SMC performs the functions of a conventional disk controller, e.g., data read/write. This state holds until the start of the next cycle.

7.3 Storage allocator

The Storage Allocator is charged with managing the storage available for digitized voice in the SMC community. This storage is addressed by specifying the SMC, Disk Transport, and segment. The number of SMCs and Disk Transports varies with office configuration; segment numbers are a property of disk geometry. The basic strategy of monitoring the idle/busy status of segments and providing rapid allocation/deallocation of resources employs a combination of main-memory-resident segment address lists and disk resident maps. The memory lists provide for normal, rapid resource allocation/deallocation, while the disk maps maintain the idle/busy status of all resources and provide a permanent record of disk configuration.

VIII. SYSTEM MEASUREMENTS

Since 1A VSS is a switching-type entity, it requires traffic and billing data and interfaces to both the Engineering and Administrative Data Acquisition System (EADAS) and the Automatic Message Accounting Recording Center (AMARC) to make these data available to the BOC. However, 1A VSS provides an entirely new class of vertical services, CCS II, without precedent. Hence, little information exists upon which to gauge, for example, customer response to the new services or traffic engineering rules.

A single, unified data collection mechanism is provided to meet data collection requirements. For each system activity of interest a unique action is defined with regard to data content, reasons for collection, method of collection, and intended uses. Because of the different uses

of the data, two primary collection methods are provided: the peg count, typically used for traffic actions which appear on the quarter hourly reports, and the transaction file, in which the action, and associated parameters, are written into a disk file.

IX. SECURITY AND RELIABILITY

A major concern during the design of 1A VSS has been the protection of customers' messages. The two aspects of this concern are that a message should be returned only to the correct customer and that messages should not be lost. To guarantee the correctness of delivery, the identity of the customer owning the message is stored in the control portion of each message segment. During playback, the identity of the requesting customer is passed to the SMC and each message segment is validated before being played. During recording, each message segment is checked before being written to verify that its current owner is the Storage Allocator. These checks protect against several types of failures that might cause a message or a message portion to be played to the wrong customer.

Messages are protected against loss by replicating each one. System announcements and voice fragments are replicated on each SMC and are accessed via translators. This is done both for reliability and availability. Because of the lower access rates, customer messages are duplicated for reliability. To accomplish this, the following deferred duplication scheme is used. As the customer speaks the message, one copy is recorded in real time. Upon completion of this recording, a duplicate command identifying this newly recorded message is placed on a queue. A background program serves this queue as processor and SMC time permits, and records the second copy of the message. To provide the desired quality of service, a system parameter specifies the maximum tolerable elapsed time to duplication. If this value is exceeded, duplication takes precedence over other new work, until the desired level of service is restored.

IX. SUMMARY

CCS II services are provided jointly by software and hardware enhancements in No. 1/1A ESS and by a new voice processing system, 1A VSS. The software for these services has been provided in such a way as to provide economical service and to permit straightforward expansion to new voice services in the future. The 1A VSS software contains the basic voice control software functions needed for many types of voice services. These building block functions permit expansion of CCS II to meet marketing requirements.

REFERENCES

1. D. P. Worrall, "1A, Voice Storage System: New Custom Calling Services," B.S.T.J., this issue.
2. R. P. Almquist et al., "Software Protection in No. 1 ESS," Proc. Int. Switching Symposium, 1972, pp. 565–9.