INTERNATIONAL OSPS DATABASE SERVICES

Bently A. Crane, John C. Dalby, Kamla Garg, Michael H. Poepping, and Charles E. W. Ward

Bently A. Crane, John C. Dalby, Kamla Garg, Michael H. Poepping, *and* ***Charles E. W. Ward*** *are members of the OSPS International Planning and Development Department at AT&T Bell Laboratories, Naperville, Illinois. Mr. Crane, who joined the company in 1957, is a distinguished member of technical staff involved in software development for international OSPS features. He holds B.S. and M.S. degrees in physics from the University of Michigan, Ann Arbor, and an M.S. in electrical engineering from New York University, New York. Mr. Dalby joined AT&T in 1970, and is a department head responsible for delivering operator services on the 5ESS switch for the international market. He has a B.S. in applied mathematics and an M.S.E. in computer information and control engineering, both from the University of Michigan, Ann*
(continued on page 50)

AT&T's commitment to expanding its switching market base beyond North America is being realized through the development of 5ESS® switch operator services position system (OSPS) features to meet the widely varied needs of operator services customers in the international market. While many features can be provided through adaptation or enhancement of existing OSPS features, many others are unique to the international marketplace. This article describes four features developed for the international OSPS product to provide operators with specialized database access capabilities.

Introduction

Many new OSPS features have been developed to meet the special challenges of the international operator services market. These features have been built on the same call processing architecture as the North American version of OSPS. This article describes how OSPS developers met the technical challenges of satisfying customer requirements for several of the new features that use databases.

In order to provide service in limited-facilities networks, international telephone offices must make it possible to book calls for later completion, and retrieve information on previously completed calls for quoting charges or reconnecting calls. Interswitch booked call and completed call data retrieval can be provided by two paired 5ESS switches. Operators also need access to online call-handling information to help to complete calls. To meet these needs, new call processing procedures and specialized databases are needed.

In the North American market, the network databases for these needs tend to be dedicated nodes in an extensive intelligent network. Access to this network is through common-channel signaling from an array of call-handling switches. However, current OSPS applications in international networks tend to be self-contained. Thus, external database systems with appropriate reliability requirements for operator services applications are not cost-effective. The new database features developed to meet the needs of international telephone service providers use databases contained within the 5ESS switch.

Booked Calls and Call Booking

In the international telephone network, and in the national networks of some countries, often it is not possible for callers to complete calls on demand because of limited network facilities or other factors such as time of day differences, limited phone penetration in destination countries, or the inability of the national network to determine the calling number for billing purposes.

The OSPS booked calls feature allows operators to store call information in the 5ESS switch for future call completion. Callers can request a specific time when call setup is to be attempted, or they can ask to have the call attempt made "as soon as possible," on a first-in, first-out basis.

In a typical booking call, the caller gives the number and name of the called party, a requested setup time, and perhaps the desired billing method. The operator enters this data using the operator terminal, and stores the call in the switch. At the specified time, the OSPS distributes the call to an operator, who calls the original caller and called party, connects them, and releases the call from the operator position.

Operators may retrieve calls before their reservation times to change or add to the stored information, attempt an early setup (if network conditions permit), or cancel the booked call if the original caller requests it.

Booked Call Storage. Basic to the design of the booked calls feature was the location of the booked call database. Use of a processor external to the 5ESS switch was rejected as too expensive and beyond response time requirements. A 5ESS switch with OSPS may have multiple position switching modules (PSMs)[1], and operators on different PSMs may book, retrieve, and/or complete a particular booked call. Because the booked calls feature must allow access to the stored calls by *all* OSPS operators, the designers decided to locate the database in the administrative module (AM). This decision also minimizes real-time and memory usage impact in the PSM, resulting in a lower system cost.

Booked call data is stored in AM main memory to provide rapid real-time response to operator storage or retrieval requests. However, over 250 bytes of data are used for each stored call, and customer requirements dictated a capacity of 5,000 stored calls. This storage requirement in excess of 1 megabyte (Mbyte) could not be accommodated by the operational kernel process (OKP) in the AM. Also, OKP does not provide facilities for disk backup of the stored call data.

Creating a new AM kernel process—the booked call kernel process (BCKP)—overcame these limitations. BCKP communicates with OKP via an RTR message interface. (RTR is the UNIX® system-based operating system of the AT&T 3B20D processor.) Booked call software in the OKP communicates with operator actions software in each PSM via switch operating system messages. It interfaces with the OSPS automatic call distributor (ACD) to distribute the booked call records to operators. The BCKP design provides for disk backup of stored call data to meet the feature's reliability requirements. Figure 1 illustrates the booked calls feature architecture.

Call data stored in the booked calls database is organized by *destination code*—typically country codes in international gateway applications—or *number zone codes* for national toll applications. For each destination code, two types of storage are provided:

- Calls requested for specific times (i.e., wanted-time calls) are ordered by their reservation times. Wanted-time calls are grouped by reservation time into 5-minute "buckets." Telephone service providers can specify the maximum number of wanted-time calls accepted by the system for storage in each bucket.
- Calls to be completed as soon as possible (i.e., deferred calls) are held in a first-in, first-out queue.

Figure 2 illustrates the database linked-list structure of wanted-time and deferred calls for a single destination code.

Operator Interface. Software for the booked calls feature located in the PSM allows the operator to enter call data. This data may vary in type and completeness

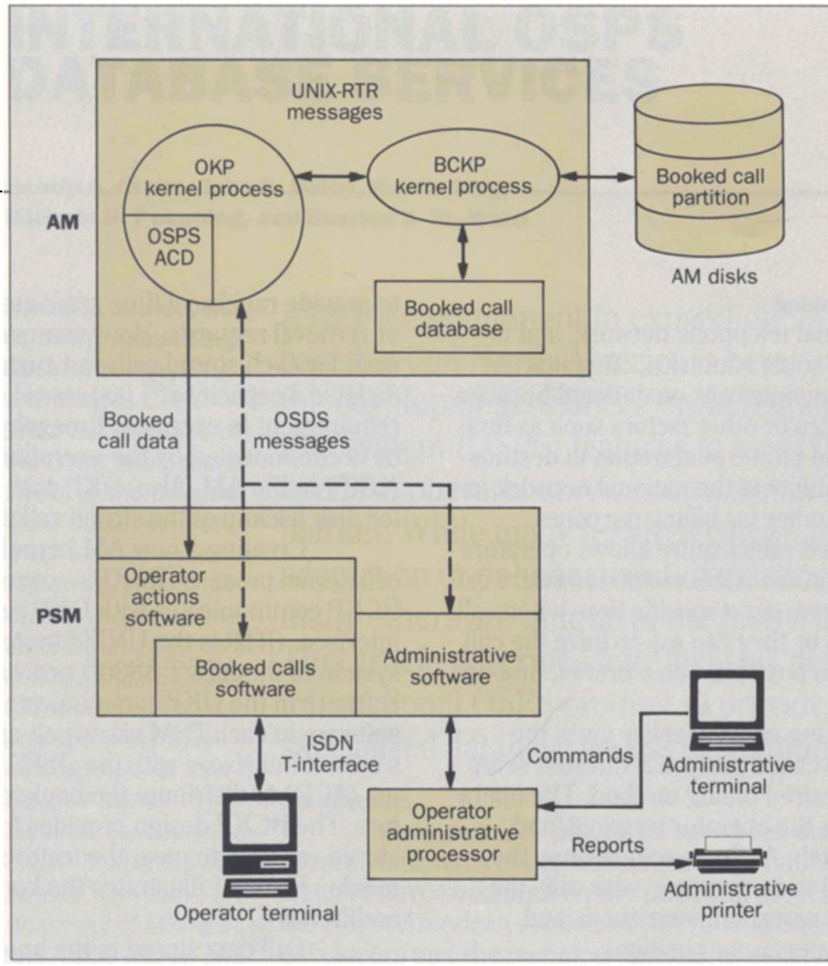

Figure 1. Booked calls feature architecture.

from call to call. The feature also allows an operator to use a subset of the caller-supplied data to retrieve the data quickly for a particular call at a later time.

The operator interface for the booked calls feature is the terminal display presented to the operator at initial position seizure, i.e., when an operator position is connected to an incoming call, thereby effectively blocking other calls until the initial call is released.[2] This display is designed so operators can treat it as a call ticket, similar to the paper tickets traditionally used by cordboard booking operators. The operator can enter call data into this "ticket" in the following format fields:

- Calling and called numbers
- Credit card number
- Third number
- Calling and called customer names
- Reservation time.

In addition, there is a free-form details field capable of holding up to 120 characters of data. A ticket number is supplied automatically to permit future reference to the stored call by an operator or by the caller.

Operators may book calls with data in any of the call ticket fields. Only the calling and called numbers must be present for a call to be stored. The operator also may enter a requested reservation time for a booked call. Entering a time in this field causes the call to be stored as a wanted-time call. Leaving the field blank results in the call being stored as "deferred."

The operator interface for retrieving calls also uses the call ticket metaphor. The operator may enter up to six different retrieval keys into a blank call ticket. The called number (alone or in combination with the calling number), credit card number and/or ticket number, can be used to retrieve a booked call. Because the number of

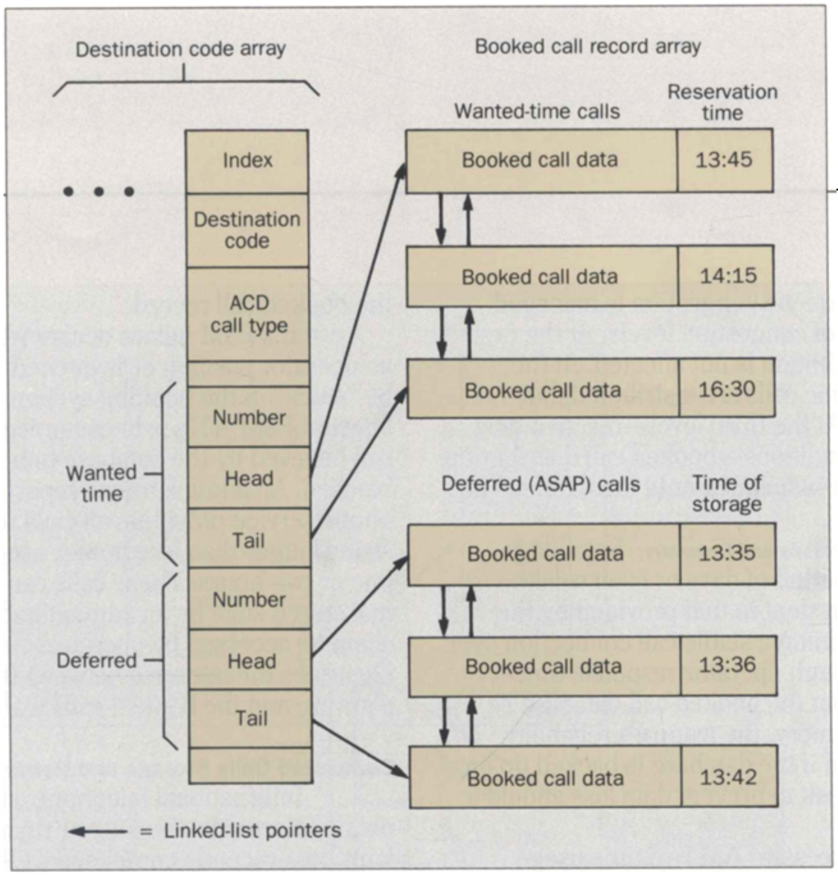

Figure 2. Booked call database structure.

stored calls is small, a sequential search algorithm is used both for its simplicity and because it provides adequate response time. When a retrieval request is handled, the first stored call encountered that matches the retrieval attributes is returned to the operator. If this is not the desired call, the operator can depress the RETRIEVE key again. The search will be continued until another match is found or until all stored calls for the destination code associated with the specified called number have been examined.

Booked Call Distribution. The booked calls feature interacts with the OSPS ACD to distribute booked calls to operators for completion. The designers addressed two challenges:

- Intermixing the distribution of wanted-time and deferred calls while respecting the designated reservation times
- Providing equitable servicing of calls to all destinations.

A round-robin distribution algorithm is used. That is, each destination code in the booked call database is visited in rotation to determine if there are calls for that destination that can be distributed. Calls from more than one destination code can be distributed to the same ACD queue. When calls are distributed, wanted-time calls with expired reservation times have priority over deferred calls. Telephone service providers can set a relative service parameter to limit how many calls may be distributed from each destination code before the next one is visited. This parameter can respond to traffic or other conditions (e.g., a limited number of trunks to a destination) by controlling the relative numbers of calls to each destination that are distributed for setup. This procedure also prevents calls to infrequently-used destinations from being blocked by heavy traffic to other destinations that use the same ACD queue. Automatic distribution of calls to a particular destination code can be suspended by assigning the value 0 to the relative service parameter.

The booked calls feature cooperates with the OSPS ACD to control overload. If congestion occurs for a particular ACD queue, queue congestion is managed through a progression of congestion levels. At the first level, booked call distribution is not affected. At the second level, wanted-time calls are distributed, but deferred calls are not. At the third level—reached only in extreme overload conditions—booked call distribution to the affected queue is suspended until the overload is relieved.

Reliability, Concurrency and Recovery. Retaining booked call data for a period of days or even weeks requires reliability equivalent to that provided by the 5ESS switch for maintaining a stable call connection over the same period. Although operator response time requirements dictate that the booked call database be located in AM main memory, the feature's reliability requirement is achieved if the database is backed up on the 5ESS switch hard disk to prevent data loss should a system failure occur.

To avoid the excessive AM real-time usage required by frequent disk writes, the entire booked call database is backed up to disk only once per hour. A complete backup is also performed automatically when the disk drives become available after an out-of-service period. The data is written to a raw disk partition, minimizing real-time consumption. A journal of changes made to the database after a full backup is built up in one of two journal buffers in initialization-protected memory in the AM. The buffer is written to the disk when it is full, or if one hour has elapsed since the last backup.

If the booked call database must be restored from disk, the in-memory database is initialized. Then, the current backup copy of the database and the update journal are read from disk and applied to the in-memory database. Finally, if either journal buffer contains valid data not previously written to disk, these changes are applied.

A booked call is purged from the database only when it either has been canceled or connected by an operator. More than one operator may retrieve a stored call, but only the *first* operator to access it can update the booked call record.

If a PSM failure occurs while a booked call is at an operator position or is queued in the ACD, the call will be "stuck" in the booking system after recovery of the affected PSM. This is because regular distribution of a call believed by the database to be in such a state is blocked. An administrative report identifies to the telephone service providers all calls remaining in these states longer than a customer-specified interval, normally one or two hours. These calls can be returned to the normal stored state by an administrative command, and can again be accessed by operators for handling. Figure 1 illustrates the interface between the OSPS administrative software and the booked calls feature.

Completed Calls Storage and Retrieval

International telephone customers may want to obtain information about—or report service quality problems on—recently completed calls made through the national or international network. For a service quality problem, call information may need to be recorded for later analysis, to credit the customer for some or all charges, or to reconnect the call. The completed call storage and retrieval (CCSR) feature provides a database for locating the call record of each completed call. It also has facilities for OSPS operators to retrieve and revise completed call information and reconnect calls.

Determining the location of the CCSR database in the distributed processor architecture of the 5ESS switch was the primary design decision for the CCSR feature. The size and content of a completed call record (CCR) ranges from 48 to 232 bytes, depending on the call type (e.g., direct dialed or operator assisted), and storage capacity for the most recent 100,000 CCRs is required. A disk-based CCR database economically meets this storage capacity requirement. However, choosing a disk-based database increased the challenge to devise an access method that provides adequate storage and retrieval performance for a mixture of variable length CCRs.

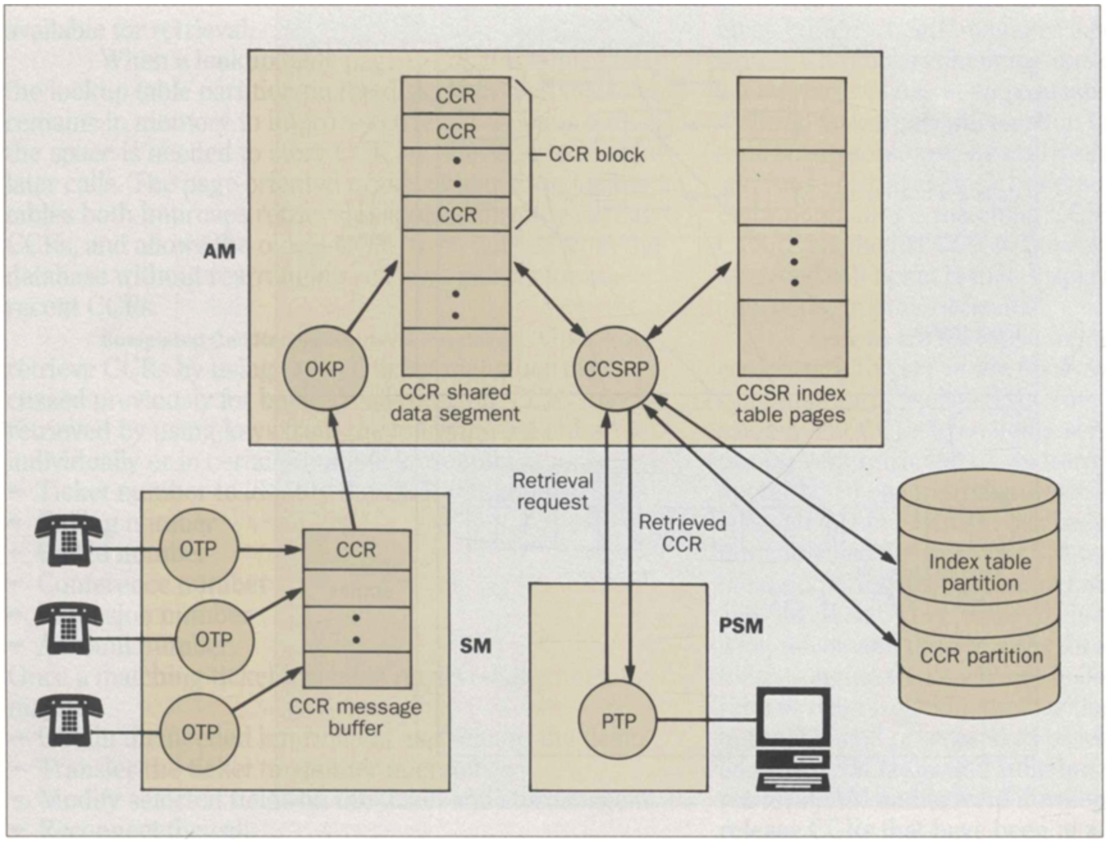

Figure 3. CCSR feature architecture.

Generation and Storage of CCRs. Figure 3 illustrates the system architecture for generating and storing CCRs. At the end of a call, the originating terminal process (OTP) populates the CCR structure for that call type and copies the record into a message buffer. The CCRs in the buffer are sent to the AM either periodically or when the buffer is full.

A process running under OKP receives the CCRs from the switching modules (SMs). This is because OKP has facilities for efficiently handling real-time, critical, call-processing messages. But because OKP does not support AM disk access, the CCSR process (CCSRP) is designed to interface with OKP to co-manage a shared data segment in AM memory. Through this data segment, CCSRP collects CCRs sent to OKP from the SMs. Like BCKP, CCSRP communicates with the RTR disk driver to write CCRs on the AM disk. The shared data segment is partitioned into blocks into which OKP loads the CCRs as they arrive from the SMs. When a block is filled, CCSRP is notified and updates the CCSR database.

Completed Calls Data Base. Since CCRs are typically retrieved within minutes of a call being completed, the database is designed to provide better retrieval response for more recent calls. The database consists of the CCRs

Figure 4. Conceptual view of CCR lookup table page.

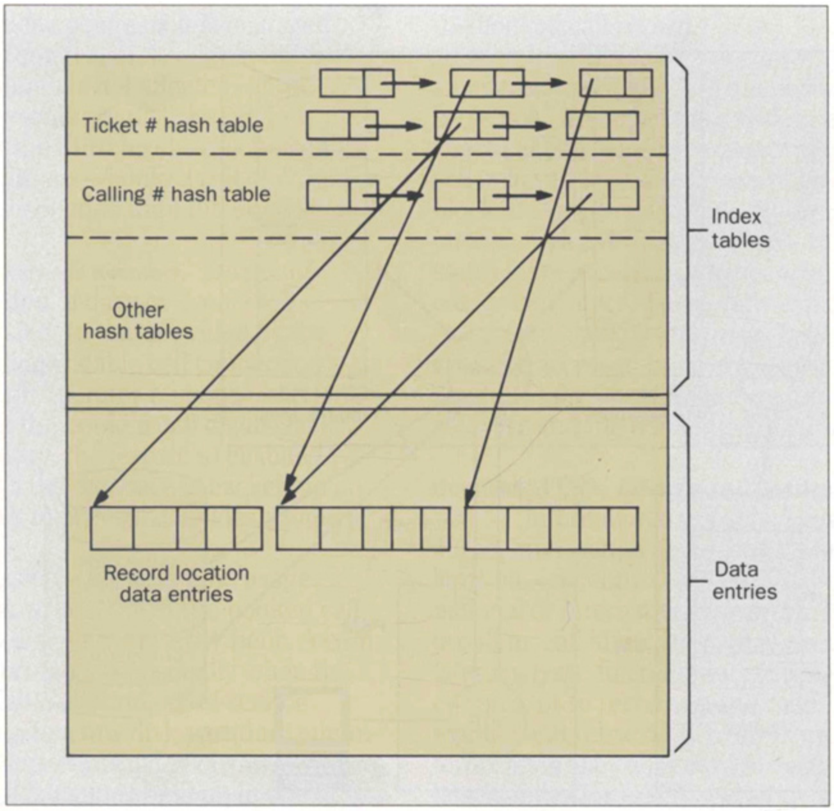

and a multikey lookup table that facilitates retrieval. CCR storage is split between the shared data segment and the CCR disk partition. After CCSRP has been notified that a given number of blocks in the shared data segment are full, it copies their contents to disk. However, the CCRs in a given block remain in memory for faster retrieval until the block is needed to store CCRs for later calls.

The CCR lookup table is an array of pages, each containing an index table area and a CCR location information area. Figure 4 illustrates a CCR lookup table page. Lookup keys are stored in one of five hash tables, corresponding to various key types. Each hash table contains hash buckets, the contents of which are stored in a linked list.

When a CCR is to be stored, its location information is written in the data area of the current page. Each lookup key is then hashed into a hash bucket for that key type. The key value is inserted in the linked list for that hash bucket with a pointer to the CCR location information. Keys are inserted at the head of the linked list to provide faster retrieval for more recent CCRs. After CCSRP is notified that a block of CCRs in the shared data segment is filled, it stores each CCR's lookup keys and location information in the lookup table. The CCRs are then available for retrieval.

When a lookup table page is full, it is copied into the lookup table partition on the disk. However, it also remains in memory to improve retrieval response until the space is needed to store CCR location information for later calls. The page-oriented organization of the lookup tables both improves retrieval response for more recent CCRs, and allows the oldest CCRs to be purged from the database without rearranging retrieval indices for more recent CCRs.

Completed Call Record Retrieval Interfaces. Operators retrieve CCRs by using the call ticket metaphor discussed previously for booked call retrieval. CCRs may be retrieved by using keys from the following list either individually or in certain multiple key combinations:

- Ticket number to identify the CCR within the system
- Calling number
- Called number
- Conference number
- Extension number
- Account number.

Once a matching ticket has been retrieved, the operator may:

- Obtain the needed information and release the ticket
- Transfer the ticket to another operator
- Modify selected fields on the ticket and store it again
- Reconnect the call.

The next match may also be requested if the ticket retrieved is not the desired one. Concurrent retrieval of the same CCR is allowed, though only the first operator has update permission.

After the operator enters the key to select a CCR, the position terminal process (PTP) sends the retrieval request to CCSRP. CCSRP searches the lookup table pages in reverse chronological order, first reading pages from memory, then from disk, until it finds a match. The block containing the matching CCR is then read from disk (or the shared data segment, if the CCR block is still in memory) into an available CCR retrieval block buffer. CCSRP manages a pool of these buffers to store CCR blocks containing one or more CCRs for which retrieval has been requested. It also maintains retrieval task status information to manage concurrency, transfer ownership of operator read/update permission, and reuse retrieval block buffers in the pool. After the block containing a matching CCR is in AM memory, CCSRP sends that CCR to the requesting PTP. The retrieved call ticket is then displayed in the retrieval window of the operator terminal.

Integrity and Recovery. In the loosely coupled processor architecture of the 5ESS switch, a system failure can temporarily isolate SMs from the AM. This failure can prevent CCRs from being sent to the AM, and also can prevent retrieved CCRs from being released by an operator. To ensure that new CCRs are not lost if the SM is isolated, a mechanism has been designed to load them into initialization-protected SM memory. (Isolation means a period loss of communication between the AM and SM. If an SM is "isolated" in that links to the AM went down, and the SM goes through some software initialization, the CCRs will not be lost.) When the SM returns from isolation, the backlog of CCRs will be sent to the AM and processed as previously described. To address PSM faults and isolation while CCRs are out for retrieval, AM audits were developed to check for and release CCRs that have been in a "retrieved" state for an unacceptably long period.

Interswitch Retrieval

Some international configurations consist of two 5ESS switches connected so calls can be completed or booked by either switch, and call inquiries can be handled by operators at either switch. The interswitch retrieval feature enables an OSPS operator to retrieve completed or booked call records stored in another switch—as well as from the operator's host switch—for a dual-switch configuration. This feature is built on the single-switch versions of the booked call and completed

Figure 5. Interswitch retrieval architecture.

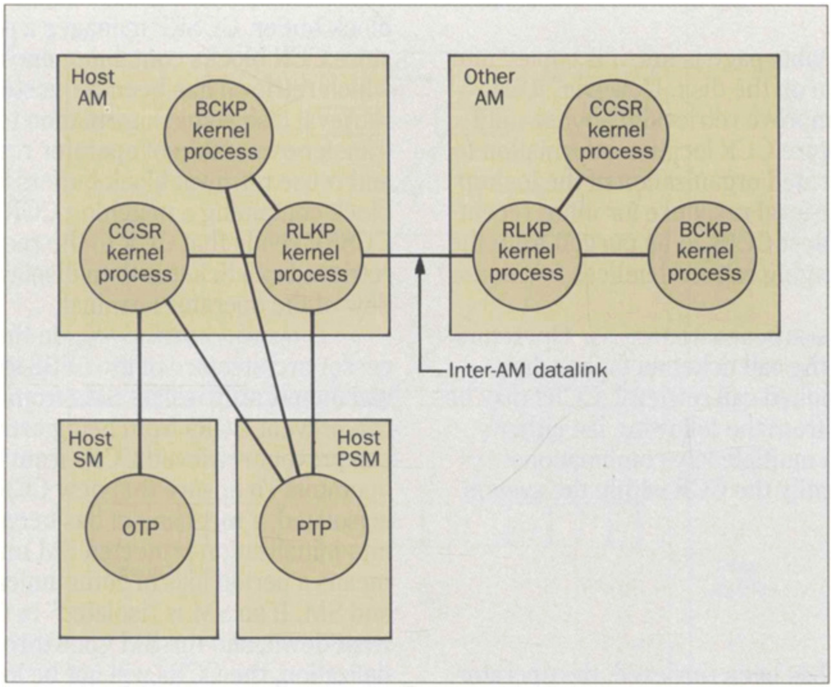

call retrieval features.

When an operator initiates a search for a completed or booked call record, the system searches the database in both the operator's host switch and the paired 5ESS switch. If the search is successful in either switch, the operator sees a display of the call ticket. Whether one or both switches are searched is transparent to the operator.

Feature Architecture. Two major design decisions were made while implementing the interswitch retrieval feature:

- Where to store the large retrieval databases of the two switches to allow access by both, but still meet the retrieval performance requirement
- How to connect the communications data link between the database and each of the two switches.

The distributed architecture of the 5ESS switch allowed several database configuration options to be considered. An external processor was considered, which would hold the database for both switches. This option was discarded, however, because of the added equipment and maintenance cost to the customer, additional development cost, and reliability concerns. Another alternative was to have an SM on one switch maintain a single database for both switches. Although this alternative would save the cost of an external processor, it would be a radical departure from the existing single-switch design, and would require a major development effort.

It finally was decided that each switch would continue to maintain its own database. A retrieval request at one switch would initiate searches on each switch, and requested records would be displayed at the requesting switch. Two types of connections were considered for providing communication between switches: a PSM-PSM link or an AM-AM link. Although a PSM-PSM link would lessen the real time used in the AM, it would not meet the retrieval delay requirements because of added message traffic between the CCSRP (AM) and the PSM retrieval link kernel process (RLKP). Hence, the AM-AM link was chosen because it required the least development effort and hardware cost while meeting the performance requirements.

Figure 5 shows the high-level software architecture for interswitch retrieval. Interswitch access is via two 9600 bits per second (b/s) X.25 data links between the AMs. The RLKP in each switch serves as the logical endpoint of the data link. It acts as a message forwarder between the two switches and uses standard RTR link interface software provided by the RTR input/output (I/O) driver.

Search Strategies. When call traffic is distributed between two switches, a customer's call to inquire about a previous call often may be routed to the switch not containing the call record. Thus, cost/performance tradeoffs become important in selecting the search strategy to meet the record retrieval delay criteria.

Two possible search strategies have been considered: sequential search (SS) and concurrent search (CS).

- *Sequential search.* SS involves searching the database on the operator's host switch before searching the database on the other switch. If the requested record resides on the other switch, a no-match search of the host switch would have to be completed before searching the database on the other switch. This strategy would significantly increase retrieval delay.
- *Concurrent search.* In the CS approach, both the host and other switch databases are searched in parallel. This strategy decreases the average retrieval time, but wastes real-time in the AM of the switch without the record.

Because the CCR database is much bigger than the booked call record (BCR) database, retrieval delay analysis was done for the CCR database using both search strategies. Estimated average CCR retrieval delay using SS was 2.35 seconds, with a maximum of 8.5 seconds. This is well above the 2-second average and 4-second maximum requirements. Using CS, the estimated average delay was 1.15 seconds, with 4.15 seconds maximum. AM real-time usage estimates were done for single switch (SS) and concurrent search. CS will use considerably more AM real time (on two switches) because it will search the whole database on the switch that does not have the record. The CS strategy was chosen for CCSR because the level of increased AM real time usage on one switch using CS is not significant, and CS meets the retrieval delay requirements.

The developers' initial expectation was that using the same CS strategy for booked call retrieval would save development effort and dramatically reduce the booked call retrieval time. However, neither benefit could be realized. A booked call record does not fit into a single interprocessor message. Therefore, storage and retrieval of BCRs requires a two-message interface between BCKP and PTP. Using concurrent search would complicate the algorithm significantly, requiring more development effort. The booked call database is small and is maintained entirely in the AM main memory. Therefore, the time required to search for a BCR is an insignificant part of the retrieval delay as compared to the delay added by the data link. Thus, the SS strategy was selected for BCR retrieval because SS met the retrieval delay requirements and minimized software complexity.

Recovery. The most challenging area of the interswitch retrieval feature has been the recovery of the CCR/BCR database when a CCR or BCR is retrieved from the other switch, or communication with the other switch is lost. RLKP never overwrites the send buffer until it receives an acknowledgment that the contents have been received. While RLKP is awaiting this acknowledgment, all new messages it receives to be sent on the data link are queued in a protected application segment (PAS) of AM memory. A data link failure can take up to five minutes to detect and report. In the meantime, many messages may get queued in the PAS awaiting acknowledgment, making recovery of the other switch's database very difficult.

If the data link is down for an extended period, special actions are needed for the BCRs that were retrieved from the remote switch. This is because booked calls cannot stay in the RETRIEVED state indefinitely. The PTP expects an acknowledgement of every message it sends to BCKP, indicating that the BCR has been put into the STORED state. Therefore, if the PTP times out while awaiting a response to a message that resulted in storing a BCR at the remote switch, it does the following:

- Assumes the data link is down
- Re-books the call on its own switch
- Purges the booked call on the other switch as soon as the data link comes up.

The recovery action for the CCR database is different. If the modifiable copy of a CCR from the other switch is used for reconnection, or if it gets modified, a new CCR record will be generated and stored in the operator's host switch, and a message will be sent to the other switch telling it to mark the old CCR as unretrievable (i.e., saved in PAS by RLKP). This is a normal operation and happens every time a modification or reconnection is done, not just when the link is down. But if the link is down, that message cannot be transmitted. The CCR must stay in the retrieved state until the link comes back up and the message to mark it unretrievable can be transmitted. The switch that stores the old retrieved CCR will know whether it was modified or reconnected only when the link comes up, so the switch keeps the entry in the "retrieved" state.

An AM audit runs periodically to remove entries of retrieval requests that were made more than an hour ago, regardless of the switch from which the request originated. However, this audit will not remove entries for retrieval requests by the other switch if the link is down, or if it has been up only a short time (to allow RLKP to clear its PAS queue of outstanding messages). If the link is down for an unacceptably long period, records that had been retrieved by the other switch can be overwritten by new records because of limited space on the disk. Entries for these overwritten records will stay in the database indefinitely because the AM audit would not remove them. A new audit was designed to run only when the link is down and will remove these superfluous entries. This is done when the CCR lookup table pages are overwritten with new records.

Visible Instruction Table

The visible instruction table (VIT) feature provides a switch-resident database that can be populated with information retrievable by OSPS operators for use in handling calls. This information is organized as free-formatted text pages that, in international OSPS gateway applications, typically contain dialing codes and rates for calls to other countries. An OSPS operator retrieves a read-only copy of a VIT page by entering the appropriate retrieval keywords, and the VIT page is displayed in a 13-line by 39-character area of the operator's screen.

The maximum capacity of the VIT database ranges from 1,000 to 2,000 VIT pages, depending on the average number of characters per page. The average delay in retrieving a VIT page is under 2 seconds. Adding, changing, and deleting VIT pages in the VIT database is done from an OSPS administrative terminal using the Recent Change/Verify facility of the 5ESS switch.

VIT Display and Keys. A VIT page is retrieved from the VIT database by specifying the values of two retrieval keywords: *VIT code* and *VIT page number*. As many as 300 different VIT codes can be defined. The allowable ranges are 1 to 1999, corresponding to the range of country codes defined by the International Telegraph and Telephone Consultative Committee (CCITT), and A to Z for retrieving information that is not country-specific. As many as 32 VIT pages can be defined for each VIT code, numbered consecutively starting at 1.

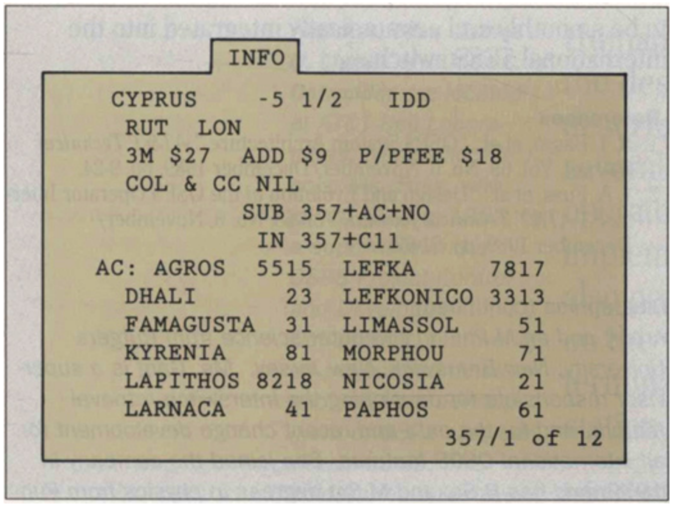

Figure 6. VIT page displayed on operator terminal screen.

An operator retrieves a VIT page by pressing the INFO key, entering a VIT code and VIT page number, and pressing ENTER. After the requested page is displayed, other pages for the same VIT code can be retrieved in ascending or descending numerical order (with wraparound) by repeatedly pressing the PG FWD or PG BK key. Only one VIT page can be displayed at a time, and it is cleared from the operator's screen upon releasing the call that was active when it was retrieved.

Several shortcuts have been provided for operator convenience. One enables a VIT code to be given an alternative three-letter "alias" that may be easier to remember and enter. Another consists of having VIT page 1 be the default if no VIT page number is entered, and having the country destination code for the currently active call be the default VIT code if no VIT code is entered. A third allows entering the asterisk character (*) to indicate the VIT code of the currently displayed VIT page.

Figure 6 shows a typical VIT page as it would appear in the upper lefthand corner of the operator's screen. The top 12 lines of the page, containing the requested information, were retrieved from the VIT database. Line 13, which identifies this page as being page 1 of 12 pages for VIT code 357, was composed by the retrieval software. This line also is used for data entry and for indicators that appear when special keys used by the VIT feature are pressed. The INFO tab at the top of the page remains visible when the page is hidden by the displays for other features that share this screen area. A hidden VIT page is displayed when any of the VIT keys is pressed.

VIT Database. One of the main customer requirements for the VIT feature was that it must provide up to 512 kbytes of free-formatted text pages. Several factors have led to the decision to store the VIT data in the switch as office-dependent data (ODD) rather than in an external database. In international OSPS applications, it may not be economical for the customer to buy and support a separate external operator-information database to serve only one or two switches. Moreover, storing the VIT pages as ODD makes use of existing administrative tools for adding, changing, and deleting VIT pages. Finally, internal storage allows a VIT page to be handled by the switch and the operator's terminal in the same way as call data, thereby allowing both to be displayed at the same time.

The architecture of the distributed 5ESS switch provided several alternatives for locating the VIT database within the switch. From a memory-conservation viewpoint, the best location would have been in the AM, where a single 512-kbyte VIT database could serve the operators on all PSMs. However, the resulting high AM-PSM message traffic and AM real-time usage ruled out this alternative. Another alternative—distributing the database among the PSMs, with each PSM containing a fraction of the total database—would have caused high PSM-PSM message traffic as well as engineering and administrative complications. The chosen alternative was to replicate the entire 512-kbyte VIT database in each PSM. Although it uses more total memory than the other alternatives, this scheme has the advantages of:

- Requiring no intermodule messages
- Providing the fastest retrieval response
- Being easier to engineer and administrate.

The VIT database stores the VIT pages as text strings of 64 characters each, where as many as eight of these strings may be needed to represent a single page. In practice, however, fewer than eight text strings are usually needed because of text compression techniques that represent trailing blank spaces and trailing blank lines as single, specially defined characters that are stripped when a retrieved page is formatted for display.

Other VIT Uses. The VIT database, although designed to provide operators with call handling information, can be used for other operator information applications. For example, an office bulletin (OB) feature, built on the VIT feature, has been developed to enable OSPS administrators to post messages to designated OSPS operators.

An administrator creates an OB by entering the message text into a VIT page and, with a separate command, specifying the operator serving teams and/or operator services centers of the operators for whom the message is intended. An OB indicator is displayed on the screens of all those operators, specifying the VIT page that should be retrieved in order to read the message. If an operator has more than one OB pending, the indicator specifies how many are pending and the VIT page of the oldest.

Summary

This article has described several OSPS features that have required special technical solutions when developed for the international market. The flexible software and hardware architectures of the 5ESS switch and of the OSPS application have allowed these solutions to be smoothly and economically integrated into the international 5ESS switch.

References

1. R. J. Basso, et al., "OSPS System Architecture," *AT&T Technical Journal*, Vol. 68, No. 6, November/December 1989, pp. 9-24.
2. J. A. Fuss, et al., "Design and Evolution of the OSPS Operator Interface," *AT&T Technical Journal*, Vol. 68, No. 6, November/December 1989, pp. 51-62.

Biographies (continued)

Arbor; and an M.Phil. in computer science from Rutgers University, New Brunswick, New Jersey. Ms. Garg is a supervisor responsible for developing the interswitch retrieval feature, and for the data and recent change development for all international OSPS features. She joined the company in 1979, and has B.Sc. and M.Sc. degrees in physics from Punjab University, India, and an M.S. in computer science from the Illinois Institute of Technology, Chicago. Mr. Poepping is a supervisor working on test environment and development coordination for international OSPS feature development. He has a B.A. in mathematics from St. John's University, Collegeville, Minnesota, and an M.S. in computer science from North Dakota State University, Fargo. Mr. Poepping joined AT&T in 1980. Mr. Ward is a supervisor of the OSPS requirements and test group, where he is engaged in feature planning and product testing for international operator services development. He joined AT&T in 1979 with a B.S. in science-engineering from Northwestern University, Evanston, Illinois, and a Ph.D. in physics from the Massachusetts Institute of Technology, Cambridge.

(Manuscript received July 31, 1989)