AT&T Technical Journal
Vol. 64, No. 1, January 1985
Printed in U.S.A.

System 75:

Communications and Control Architecture

By L. A. BAXTER,* P. R. BERKOWITZ,* C. A. BUZZARD,†
J. J. HORENKAMP,* and F. E. WYATT*

(Manuscript received July 11, 1984)

The System 75 office communication system uses a unique communications and control architecture that provides great flexibility and a minimum of overhead for small configurations while growing smoothly to larger line sizes. A distributed communication network provides 64 kb/s connectivity for both voice and data. It consists of a pair of time division multiplexed buses and intelligent port circuits. Flexible conferencing and gain adjustment are supported as an integral part of the network. The control complex supports an operating-system-based software structure.

I. INTRODUCTION

The System 75 office communications system hardware architecture consists of a control complex and a communications network, which are connected by a pair of Time Division Multiplexed (TDM)‡ buses, as shown in Fig. 1. Part of the bandwidth of the TDM buses is used as a control channel between the control complex and the intelligent port circuits in the communications network.

* AT&T Information Systems Laboratories, an entity of AT&T Information Systems, Inc. †AT&T Information Systems Laboratories; present affiliation, Bell Communications Research, Inc.

‡ Acronyms and abbreviations used in the text are defined at the back of the *Journal.*

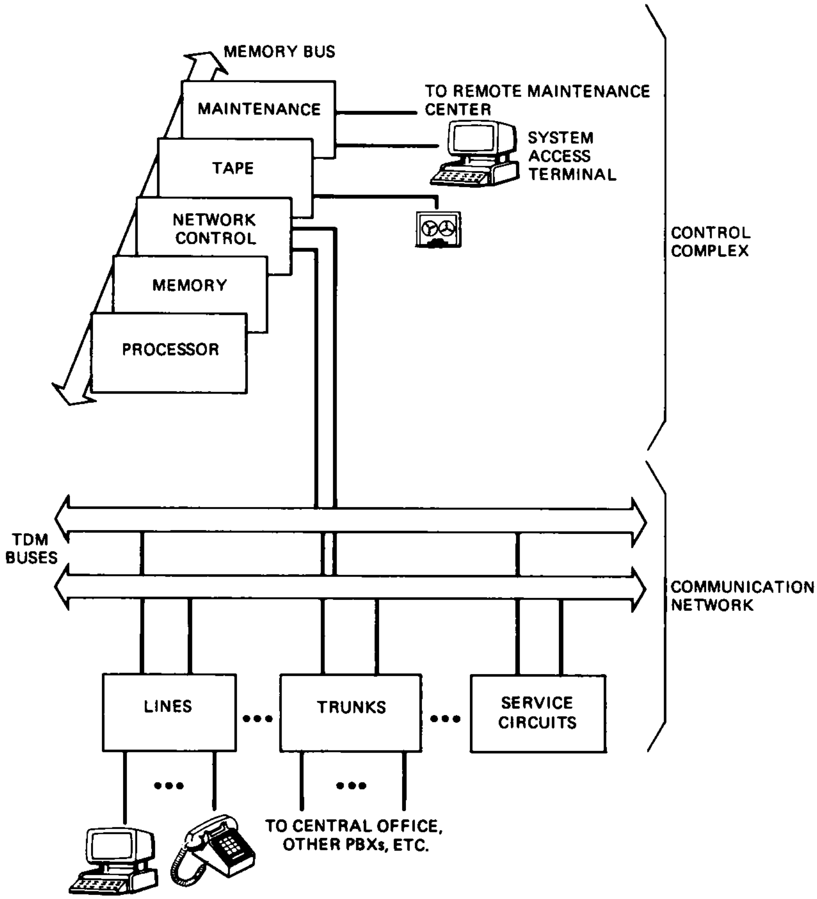

Fig. 1—System 75 communication and control architecture.

One of the main architectural features of System 75 is the distributed switching network. This was chosen to allow as much complexity as possible to be transferred to the port boards, thereby reducing the amount of common circuitry required and minimizing the getting-started cost (often referred to as the intercept cost). At the same time, aggressive use of VLSI devices in the port circuits allows the per-port cost (slope) to remain low. This combination of low intercept cost and moderate slope allows the System 75 network to remain cost-effective over a wide range of line sizes.

Another key architectural feature is the universal slot concept. All circuit pack slots in System 75 port carriers are identical (with the exception of a unique address for each slot). Every port slot has the same interfaces to the TDM buses, I/O access to outside devices, and power supplies. This freedom to plug any port circuit into any slot allows customers to flexibly configure their system, so that it is optimized to their particular needs.

This highly distributed and modular architecture has allowed System 75 to meet the following design objectives:

1. Provide a digital network which efficiently supports both voice and data communication.
2. Serve customers up to four hundred lines with a single-cabinet hardware configuration.
3. Provide support for an operating system-based software structure.
4. Maintain high-reliability operation with complete self-diagnosis and alarming.
5. Utilize a flexible architecture so that future needs may be easily accommodated, and the system may be gracefully upgraded as technology advances.
6. Provide the above functions at a competitive price.

The following sections describe the architecture in more detail.

II. COMMUNICATION NETWORK

2.1 TDM buses

System 75 has two parallel TDM buses, each of which is 8 bits wide and runs at 2.048 MHz. Functionally, the dual-bus structure is equivalent to a single 512-time-slot bus. Separating the bandwidth into two physically distinct buses has two advantages. First, it lowers the speed of each bus, which eases the timing requirements on VLSI interface devices. Second, the redundancy provided by two buses improves system reliability. If one bus fails, the architecture permits continued operation at reduced capacity on the other bus.

The buses are implemented as printed paths on the backplane. The geometry of these printed paths has been carefully designed to maintain the proper characteristic impedance. Several carriers may be daisy-chained together within a cabinet. Each bus path has a resistive termination at each end.

A novel current-source bus transceiver was designed for this application. Up to 100 port boards may be plugged into the bus in a simple party-line fashion. These transceivers have been specifically designed with low signal levels and controlled rise times to minimize radiated noise. They are designed to allow nondisruptive insertion and removal of boards with the system power on and to isolate port boards from the bus during failure conditions.

Voice signals on the buses are encoded in μ-255 PCM (Pulse Code Modulation) format[1] for domestic systems, while A-law PCM could be provided for international applications. Data signals utilize the Digital

SPECIFIC ADDRESS

MSB 7	6	5	4	3	2	1	LSB 0
0	A6	A5	A4	A3	A2	A1	A0

A0 - A6 = ANGEL ADDRESS

GROUP ADDRESS

MSB 7	6	5	4	3	2	1	LSB 0
1	G3	G2	G1	G0	T2	T1	T0

G0 - G3 = GROUP ADDRESS
T0 - T2 = SCAN/COMMAND TYPE

Fig. 2—Addressing modes of TDM bus control channel.

Communications Protocol (DCP).[2] Several time slots are reserved for tone distribution and for the control channel.

2.2 Control channel

The first five time slots on each bus are reserved for a control channel between the control complex and the port circuits. In essence, the control channel is the backbone for a network of microprocessors. It provides a communication path between the control complex and the microprocessor on each port circuit pack (commonly referred to as the angel). On each port board a custom VLSI device known as the SAKI (Sanity and Control Interface) provides address recognition, buffering, and synchronization between the angel and the five control time slots. The control channel is active on only one bus at a time. It can be moved to the other bus in the event of a bus failure.

The control channel operates strictly in a polled mode, with the network control (often referred to as the archangel) as master, and the angels as slaves. The first time slot of each frame (TS0) carries the control address, while the following four time slots (TS1 to TS4) contain control data. The archangel grants bus usage to a particular angel or group of angels by transmitting a specific address in TS0. The direction of transmission during the control data time slots (TS1 to TS4) is dependent on the message type.

Each port slot in System 75 contains seven address pins that are hard-wired to define a unique address. During initialization, the angel reads in this address and writes it to the address-detection portion of the SAKI. This seven-bit address fixes a limit of 127 angels (plus one null address) in the archangel's address space. This restriction is well above the physical limitations on the number of port circuit packs in a single cabinet.

There are two modes in which the archangel can address an angel, as shown in Fig. 2. To differentiate between these two modes, the SAKI inspects the Most Significant Bit (MSB) of the control address which appears in TS0. If the MSB = 0, then the archangel is addressing the specific angel whose address is given in the remaining seven bits. If the MSB = 1, then the archangel is addressing the group of eight angels whose address matches bits 3 through 6 of TS0. In this case, the three Least Significant Bits (LSB) indicate the type of scan or command.

2.2.1 Group addressing

The group address mode is used in two ways: to collect status information (called short-scanning) and to send certain commands to a group of eight angels simultaneously (group commands). In a short scan, each angel in the addressed group responds with a single bit of information in TS2. The bit-position assignments on the TDM bus are determined by a binary decoding of the lower three bits of the angel address. (In other words, the angel whose lower three address bits are 000 responds on TDM bus bit 0, etc.) Using short scans, the archangel can gather status information from a full complement of angels an order of magnitude faster than if each angel were polled directly. Thus, short scanning reduces the latency period before a stimulus is reported to the control complex.

The archangel obtains two types of status information via short scans. Activity scans determine which angels have messages waiting for uplink transmission. Those angels will be individually polled to collect the messages. Sanity scans collect state-of-health information from the angels.

The sanity control circuitry in the SAKI gives the control complex the ability to identify and isolate insane angels quickly. When the archangel sends a sanity scan to a group of eight angels, each SAKI in the group checks the sanity of its angel by verifying that its angel has updated a special sanity bit latch in the SAKI since the previous sanity scan. If its angel has cleared the sanity bit, the SAKI notifies the control complex of its angel's health by driving its bit low during TS2. If the angel has not cleared the sanity bit, indicating angel insanity, the SAKI sends no response in TS2, notifying the Switch Processing Element (SPE) of the angel's insanity. Simultaneously, the SAKI resets its angel, holding it idle, and forces the bus transceivers into a receive-only mode, preventing the port board from errantly transmitting onto the TDM bus. The SAKI waits for the restart instruction from the control complex before allowing the angel to begin running again.

Unlike traditional watchdog circuits where each processor must update a local timing circuit periodically to indicate sanity, System 75's sanity control gives the control complex total control over the sanity scanning rate, the number of times an insane angel is restarted, and the ability to shut down an individual angel at will.

Each SAKI also protects the control channel by monitoring transmission onto the TDM bus during the control time slots. When it detects transmission during a control time slot by anyone other than itself, it disables its angel and Network Processing Elements (NPEs) and waits for the control complex to send the restart command.

2.2.2 Specific addressing

When the control complex wants to send a message to a specific angel, or retrieve a message from an angel that gave a positive response to an activity scan, the archangel must use the individual addressing mode. In this mode, a message may span a number of frames.

Messages sent across the control channel use a well-defined format known as the Control Channel Message Set (CCMS). The CCMS provides a combination of stimulus and functional messages that are common across all types of ports. Downlink (network control to port circuit) messages allow the control complex to control the ringer and LEDs on stations; seize, release, and outpulse on trunks; set up and tear down network connections; execute various maintenance tests, etc. Uplink (port circuit to network control) messages allow the port to report state changes, such as switchhook and button pushes on stations and seizure of incoming trunks. Control channel messages are protected by a checksum, and are retransmitted in the event of an error.

2.3 Network processing element

2.3.1 Switching functions

In conventional digital switches, each port is permanently assigned a time slot on which to talk and another on which to listen. A centralized mechanism called a Time Slot Interchanger (TSI) is used to enable, reorder, and transfer time slots from talking to listening ports. Conventional TSIs require additional centralized equipment to perform gain adjustment or form conferences, features that require arithmetic processing of voice samples. However, intelligent TSIs have been designed to perform these operations.[3]

Whether intelligent or not, a centralized TSI that is sized to accommodate the full capacity of the communications system represents a cost burden on small-size customers who pay for more capacity than they need. The alternative of providing a family of TSIs optimized for several sizes entails extra development effort and complicates growth in the field.

The communications network in System 75 solves these problems. Each port board carries with it a modular piece of the network in the form of a VLSI chip, the NPE.[4] Each NPE serves four ports and is resident on each of the port boards. Two are used on each of the eight-port boards with the exception of the digital line, which uses four NPEs to switch the two information channels of each of its eight ports. The distributed network architecture and absence of a centralized TSI allows customers to buy just the right amount of network for their needs and permits smooth growth as needs expand, while the VLSI technology provides a low per-port cost even though each port has its own dedicated TSI and voice processing logic.

The NPE provides the functions of time-slot assignment for listening and talking, gain adjustment, and eight-party conferencing.[5] (System 75 actually features a six-party conference limit with the remaining two conferencing slots reserved for tones.) The NPE contains over 18,000 transistors, with half of them making up a novel memory network for control and processing functions. As an indication of its complexity, a Transistor-Transistor Logic (TTL) breadboard of this device required six 8- by 13-inch circuit boards.

2.3.2 NPE operation

Figure 3 is a diagram of the operation of one of the NPE's four channels. A network of memory arrays, the associative conference buffer,[6] is used both as a control store, written and read by the angel, and as a buffer for PCM samples from the time division bus. Memory locations are loaded by the angel with time-slot numbers for specifying

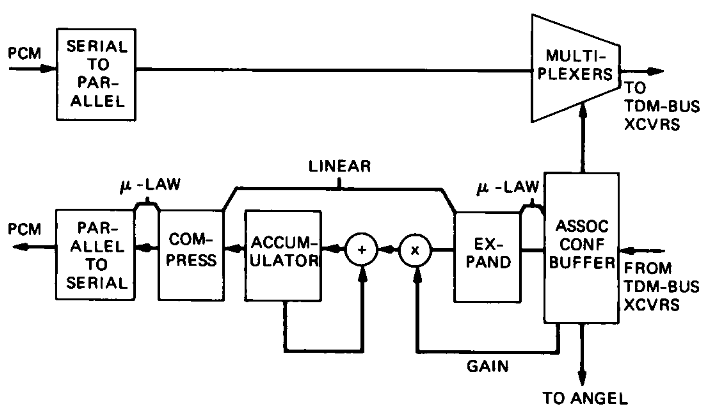

Fig. 3—Architecture of one of the network processing element's four channels.

a talking slot and up to seven listening time slots. Companion memory locations are loaded with a gain or loss value to be applied to samples from listening time slots received from the bus. A talk enable/disable bit can also be stored for the talking time slot. The locations holding the time slot number also act as a content addressable memory by comparing their content against a time-slot counter and individually controlling a sample transfer on the specified time slot. The sample transfers consist of placing a talking sample from a station onto the bus and storing up to eight listening samples from the bus in a sample buffer memory array. The sample buffer holds the samples until they are accessed with their respective gain values to form a conference sum.

An active port is usually allowed to talk on one time slot and listen to from one to seven others. An idle port uses no time slots. When multiple listen time slots are selected, their samples are converted to linear PCM, multiplied by the stored gain values, summed together in an accumulator, and then restored to μ-law for delivery to the station. Of course, data samples are passed through the NPE without any of the conference processing which would corrupt the data.

A simple two-party connection occupies two time slots: each port talks on one and listens to the other. An N-party conference uses N time slots. Tones may be broadcast on a single time slot and received by an unlimited number of ports. In System 75 the seven single-frequency components of the Dual-Tone Multifrequency (DTMF) signaling tones are broadcast continuously. Each port requiring access to a DTMF tone for dialing out forms a brief conference between the two appropriate single frequencies without interfering with other ports similarly doing so.

2.3.3 Conferencing algorithms

The forming of a gain-adjusted conference sum can be thought of as a sequence of arithmetic operations. (In the following discussion, conversions between linear and μ-law PCM are neglected, since they are required in all cases.) Each recipient of the conference sum must hear the composite of the other conferees' samples minus his or her own. The sample received by the kth member of an N-party conference is:

$$R_k = \sum_{i=1}^{N} g_{ik}T_i - g_{kk}T_k,$$

where

R_k = Receive sample for port k
T_k = Transmit sample for port k
g_{ik} = Gain coefficient from port i to port k

T_i = Transmit sample for port i
g_{kk} = Gain coefficient from port k to itself.

There are conceptually two algorithms for generating a conference sum. One method is to hold the g_{ik} constant for all k. Then:

$$R_k = \sum_{i=1}^{N} g_i T_i - g_k T_k,$$

where g_i = Transmit gain for ith port. First a conference sum of all gain-scaled transmit samples is formed, a task requiring N multiply and accumulate operations. Then the receive samples are generated by subtracting out the receiving port's scaled transmit sample, an additional N operations, for a total of $2N$ operations per conference.

The chief advantage of the "$2N$" algorithm is its efficient execution, an important property in a centralized, intelligent TSI where processing throughput may be a constraint. A disadvantage is the inflexibility it imposes on conference call transmission gains, since all conferees must listen to a given port with the same gain. Often this results in loss in excess of that required for stability or optimum intelligibility.

The second method allows individually chosen interport gains between all conferees but requires considerably more processing effort for large conferences. It builds each receive sample R_k separately:

$$R_k = \sum_{\substack{i=1 \\ i \neq k}}^{N} g_{ik} T_i.$$

Samples from the other N-1 conferees are individually multiplied by the appropriate transmission gain coefficient and added to the partial sum as it is built up. This requires N-1 multiply operations. Since the sum must be formed separately for each receiving port (a total of N times), $N \times (N - 1) = N^2 - N$ operations are needed to form the conference. The advantage of this "N-squared" algorithm is the freedom it allows in choosing transmission gain coefficients. Each party can listen to the other conferees with an individually tailored gain.

System 75 achieves an N-squared conference algorithm since each of the N NPEs in a conference perform the required N-1 operations with individually chosen gain coefficients. The necessary processing throughput is obtained as a natural benefit of the parallelism inherent in the NPE-based network.

The importance of N-squared conferencing is illustrated by a three-way call involving two telephone line ports and a central office trunk port. PBX line ports optimally require about 6 dB of loss between them to simulate losses normally encountered in the loop plant, while trunk connections should have 0-dB transmission loss between them, as shown in Fig. 4. This combination of loss relationships can only be achieved with an N-squared algorithm since a $2N$ method would subject the trunk to the same 6-dB loss that is applied between lines. The NPE implementation of this three-party conference is illustrated in Fig. 5.

2.4 Intelligent port circuits

Figure 6 is a block diagram of a generic System 75 port circuit. The

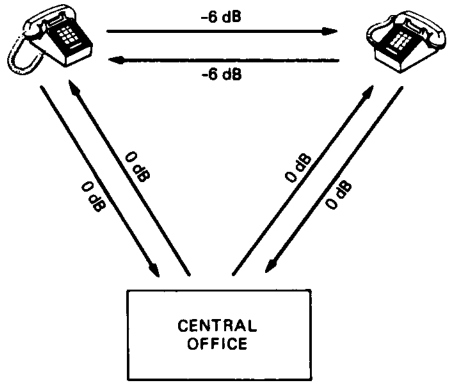

Fig. 4—Example of System 75 gain plan for three-party conference.

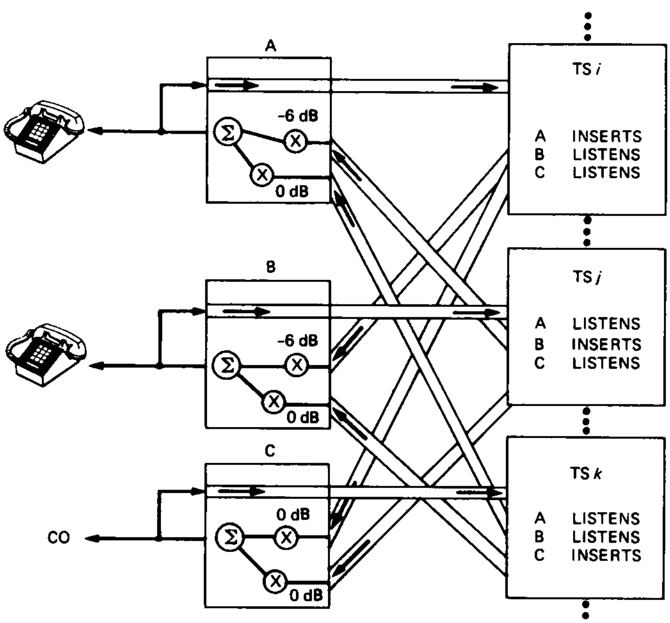

Fig. 5—System 75 implementation of three-party conference.

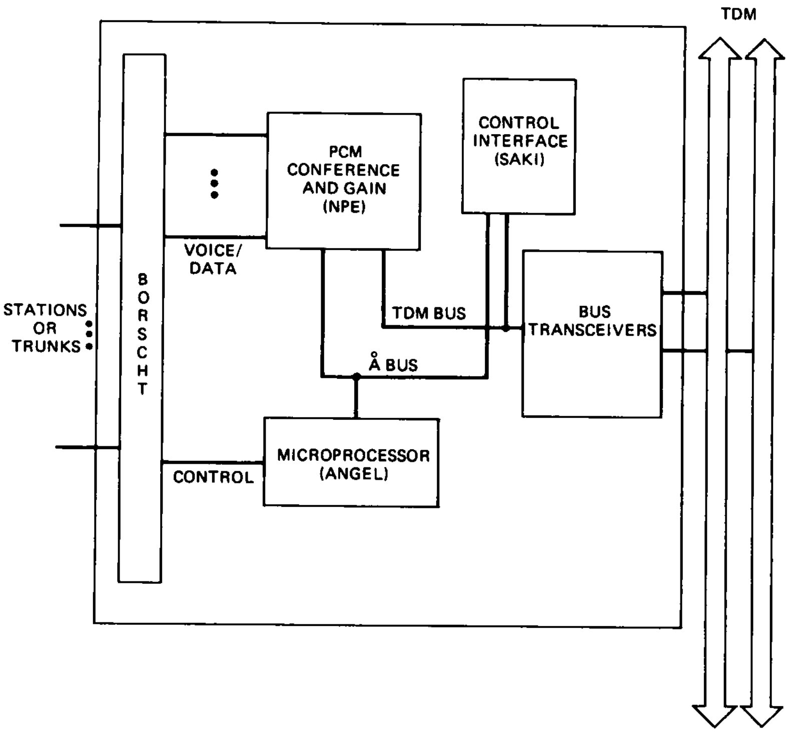

Fig. 6—Architecture of a System 75 generic port board.

circuit pack interfaces to the TDM buses via the custom bus transceivers. Time-slot information, which may be either PCM voice samples or data, is handled by the NPEs (Section 2.3). The interface to the control channel is handled by another VLSI device, the SAKI (Section 2.2).

The BORSCHT circuitry contains whatever is necessary to interface to a particular type of line or trunk. (BORSCHT is an acronym for Battery feed, Overvoltage protection, Ringing, Supervision, Codec, Hybrid, and Testing.) In general, this block of circuitry is different for every type of port circuit.

The heart of the port circuit is the on-board microprocessor, or angel. Every port board in System 75 has an angel that controls the operation of the circuit pack. The angel is implemented as a single-chip microcomputer with up to 8K bytes of firmware. The firmware is divided into two sections: a common portion, which is essentially the same for all circuit packs, and an application portion.

Both the NPE and the SAKI are operated as peripherals to the angel. When setting up a network connection, for example, the following actions occur. The control complex formulates a down-link message and sends it over the control channel. The message is received by the SAKI and passed to the angel. The angel sends an acknowledgment (via the SAKI) and examines the message. It then loads the proper time slot and gain values into the NPE so that the desired connection is established.

The distributed intelligence of the angels is a key element which makes a common control channel message set possible. The angels also play an important role in the maintenance strategy. The angel's responsibilities include:

1. Scanning the station/trunk and reporting any state changes uplink to the control complex.
2. Interpreting received (down-link) control channel messages and taking the proper action. For example, when a 'ringer-on' message is received, an analog line-circuit angel must close a relay to provide 90-volt ringing, do ring-cycle timing, etc. A digital line circuit angel, however, would format a command to the station set and send it, using the DCP signaling channel.[2]
3. Handling all short-duration timing functions. Examples include ring cycle and LED cadence timing for stations, outpulsing for trunks, and interdigit timing for DTMF receivers.
4. Performing a variety of maintenance tests. In addition to tests which are run on command from the control complex, the angel does extensive in-line error testing during the normal operation of the circuit pack. Error pegs are kept and reported uplink.

In summary, the angel provides the intelligence necessary to isolate call processing software from port-specific differences and to off-load the control complex from having to do real-time intensive port scanning functions. These are reflected in the virtualization provided in the control channel message set.

This cleanly defined, message-based control interface, coupled with the universal slot concept provides yet another benefit. It is relatively straightforward to design new types of port circuits and integrate them into System 75. Our current family of port circuits includes five types of line circuits (analog, hybrid, Multibutton Electronic Telephone [MET], digital, and data); five types of trunks (central office, direct-inward-dial, tie, auxiliary, and DS-1[7]); and four types of service circuits (tone/clock generator, tone detector, pooled modem, and speech synthesis). Most port boards provide eight port circuits. In the case of the digital line circuit, this results in 16 network appearances, since each port supports two information channels in the digital communications protocol.

Many of the station types are supported across the product family. Analog sets are supported on all AT&T Information Systems PBXs. The hybrid set was adapted from the *Merlin*™ communications system and is also supported by the *Dimension*® System 85 communication system. The MET set provides an economical migration path for customers who already own *Dimension* or *Horizon*® communications systems. The digital stations provide advanced voice/data features using the digital communications protocol and are common with System 85. This variety of ports and stations allows the system to be tailored to the specific needs of each customer in a cost-effective manner.

2.5 Digital line circuit

To illustrate the concepts previously mentioned, a particular circuit pack, the digital line circuit, is discussed in this section in more detail. Like all System 75 port circuits, the digital line circuit makes extensive use of custom VLSI devices and performs many functions in firmware rather than hardware. The digital line circuit, which terminates eight DCP lines, is an evolutionary step towards an Integrated Services Digital Network (ISDN).[8] Like the proposed ISDN interface, each DCP line provides two information channels and a separate channel for signaling, thereby supporting simultaneous integrated voice/data communication. Thus, this circuit pack supports sixteen endpoints—a density unmatched by any of the other port boards.

2.5.1 Hardware configuration

Figure 7 is a photograph of the digital line circuit. The major functional blocks are indicated on the figure. The five integrated circuits at the upper left are the bus transceivers. The SAKI device, which provides hardware support for the control channel, is at the right of the bus transceivers. Note that the SAKI, like several other devices on the board, is packaged in a 68-pin surface-mount chip carrier. (Physical design considerations are explained in more detail in Ref. 9.) To the right of the SAKI are four NPEs, which provide access to the TDM buses for the 16 information channels that this circuit pack supports. The angel microprocessor that controls the operation of the circuit pack, and its associated RAM are at the right side of the circuit pack. All of the components mentioned above are common to all System 75 port circuits, and are also shown in the generic port board diagram (Fig. 6).

The remaining circuitry (called the BORSCHT in Fig. 6) interfaces directly to the DCP lines. The bottom half of the circuit pack contains eight identical blocks of circuitry. The Digital Line Interface (DLI) device contains a complete 160-kb/s modem packaged in a 40-pin DIP (Dual In-line Package). It provides full-duplex operation over up to 5000 feet of 26-gauge cable, and includes circuitry for framing, scrambling (to reduce radiated noise), clock recovery, and automatic equal-

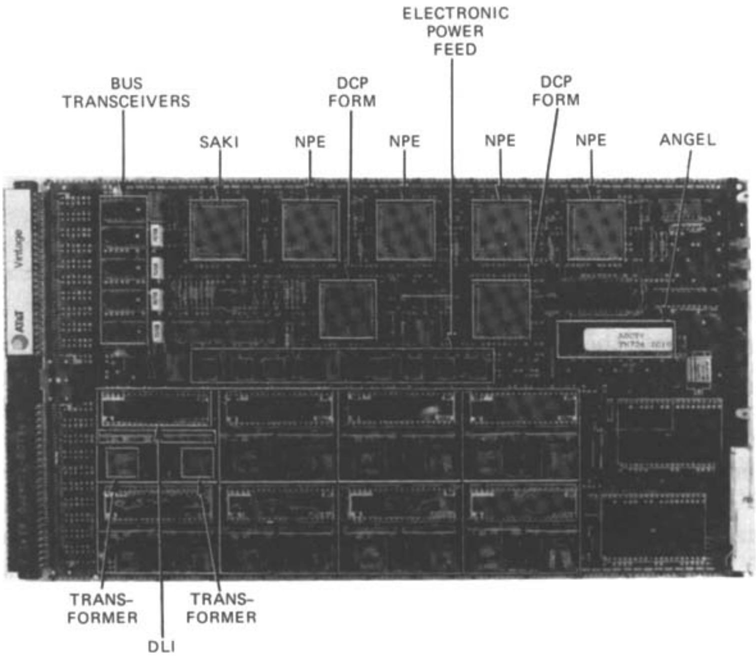

Fig. 7—Photograph of System 75 digital line circuit with major functional blocks indicated.

ization. The two SIPs (Single In-Line Packages) immediately below the DLIs contain the external resistors and capacitors needed by the DLIs. A pair of transformers provides the actual interface to the DCP lines. The use of a transformer-coupled interface has a couple of advantages: it protects the board against longitudinal surges, and it allows power to be supplied to the station over the same pairs of wire (via a technique known as "phantom powering").

The Electronic Power Feed (EPF) chips which are in the middle of the circuit pack control station power. An EPF is a microprocessor-controllable electronic circuit breaker. The EPFs automatically shut down when an overcurrent condition is detected and can also be turned on and off by the angel. In addition, the angel can read the status of each EPF and determine (1) whether it is supplying a normal amount of current to the station, (2) whether it is not supplying any current to the station (this normally means that the station is unplugged), or (3) whether it is in overcurrent mode, which indicates a fault condition in either the station or wiring.

The DCP formatters are custom integrated circuits that provide link-level hardware support for the DCP signaling channel. Since most of the DCP signaling channel protocol is implemented in firmware, further discussion of the DCP formatters will be deferred until the next section.

2.5.2 Firmware interactions

The digital line circuit angel firmware has three main functions:

1. It processes control channel messages to and from the archangel, as described in Section 2.2.
2. It translates between the control channel (CCMS) protocol and the DCP signaling protocol.
3. It performs a number of maintenance functions, such as logging and reporting transmission errors.

The digital line circuit angel firmware is built around a task dispenser known as APEX (Angel Processor Executive), while real-time I/O to the SAKI and DCP formatters is interrupt driven. Processing DCP messages will be described in more detail, since it is the most complex of the above functions.

The DCP provides an 8-kb/s signaling channel that uses a simplified High-level Data Link Control (HDLC) protocol. In particular, the framing, bit stuffing, Frame Check Sequence (FCS), and link initialization commands (SABM, DM, and UA) are identical to HDLC.[10]

The DCP formatter devices each provide link-level hardware support for four DCP links. In the uplink direction (station to PBX), the formatters provide flag detection, bit de-stuffing, and message demarcation. They generate an angel interrupt when a message byte has been assembled (approximately every millisecond during a message transfer). The angel stores the received bytes in a buffer until the formatter indicates that the complete message has been received.

The completed message is processed at base level in the next APEX task cycle. The angel calculates and verifies the FCS, checks the sequence number, and transmits an acknowledgment—called a Receive Ready (RR)—to the station. The message is then converted to CCMS format, and moved to a different buffer to await uplink transmission over the control channel.

In the downlink direction, messages received over the control channel are converted to DCP format during an APEX task cycle. This includes prepending the correct header and sequence number, and calculating and appending the correct FCS. The message is delivered to the DCP formatters one byte at a time via an interrupt-driven routine. The formatters take care of flag generation and bit stuffing. The angel retains the message in its buffer until an acknowledgment is received from the station. As in HDLC, if none is received within a specified period, the message is retransmitted a maximum of two times. If no acknowledgment is received after the third try, the angel attempts to reinitialize the link.

The two information channels on a DCP link use different logical channels for signaling. Thus, the angel must maintain 16 HDLC-like protocols simultaneously. The angel has responsibility for all link-level functions, including link initialization, sequence numbering, FCS generation and checking, acknowledgments, and retransmissions. As mentioned previously, this allows the call-processing software to maintain a uniform message-based (CCMS) interface to all types of endpoints.

2.6 Digital signal processing technology

Digital signal processing technology is used extensively on all the System 75 service circuits. The AT&T Digital Signal Processor (DSP) integrated circuit[11] is used to implement the signal processing algorithms in System 75. Its advantages include small size, high reliability, low cost, low power consumption, and the availability of numerous development tools.

Some of the many uses of the DSP within System 75 are:

1. The tone/clock circuit pack uses two DSPs to digitally generate all the various tones used by the PBX (e.g., dial tone, ringback, busy tone, intercept tone).

2. The tone detector uses DSPs to implement both Dual-Tone Multifrequency (DTMF) receivers and general-purpose tone detectors (for detecting dial tone, modem answer tone, maintenance tones, etc.) on a single circuit pack.

3. The pooled modem circuit pack contains conversion resources to convert 212A modem signals into DCP format.[2] DSPs are used to implement two 212A-compatible modems on the circuit pack. The advantages of this circuit pack over conventional modem pools include lower cost, uniform administration, better maintenance, and reduction of PBX-room clutter.

4. The speech synthesis circuit pack uses DSPs both for DTMF receivers and for generating Multiple Pulse Linear Predictive Coding (MPLPC)[12] speech samples from stored coefficients.

III. CONTROL COMPLEX

The System 75 control complex is shown in Fig. 1. The control complex is often referred to as the Switch Processing Element (SPE). It consists of a processor, memory, and I/O connected by a single-master Memory Bus (MBus). This configuration meets the cost, performance, and reliability goals for basic service and it can be expanded to support optional services.

3.1 Processor

The processor consists of a commercial 16-bit *Intel** 8086 microprocessor and a Memory Management Unit (MMU) implemented in custom gate arrays. The microprocessor and the MMU functionality were chosen to provide good performance for the largest system configurations and minimum cost for the smallest configurations. Specific constraints are:

1. To minimize equipment cost, the processor and MMU are implemented on a single circuit pack.
2. For maximum performance, most memory accesses are accomplished with only two wait states, including memory management and error correction overhead.
3. Multiple contexts and fast context switching are supported to achieve maximum operating system performance.
4. A high degree of self-checking and protection is provided for call processing applications.

Design trade-offs were made between hardware and software to meet these constraints. The result is an MMU which supports 16-bit virtual to 24-bit physical address mapping, 15 segments of up to 64K bytes each, and the following protection features:

1. Two levels of execution privilege (system and user).
2. Bounds checking on any access, with an overflow stack to aid recovery from stack exceptions.
3. Illegal instruction detection (e.g., HALT instruction).
4. Segment write-protect capability.
5. Distinction between text and stack/data segments to prevent execution of data and to provide execute-only access of text.

3.2 Memory

Because System 75 is software-intensive, the memory can have a significant impact on system cost, reliability, and performance. To meet the system design objectives, the memory uses 256K Dynamic Random Access Memory (DRAM) devices and Error Detection and Correction (EDC) logic. Each memory circuit pack provides 2M bytes organized into 22-bit words (16 data bits + 6 check bits). The EDC circuitry provides single-bit error correction and double-bit error detection and therefore dramatically improves the system's mean time to critical failure. The memory uses VLSI devices to incorporate all refresh, control, and maintenance functions on each pack, thereby eliminating any external memory control function.

---

* Trademark of Intel Corporation.

3.3 Input/output

The I/O functions are implemented with intelligent interfaces which off-load the processor and shield call processing software from real-time-critical tasks. The processor communicates with the interfaces through dual-port memories on the MBus.

3.3.1 Network control

As previously discussed, the network control circuit pack provides the bridge between the control complex and the communication network. It is the master of the TDM bus control channel and, in addition, provides a time-of-day clock with battery holdover, a system clock failure detector, and four switched data channels used for dial-up maintenance/administration and printer output.

3.3.2 Tape interface

The tape interface circuit pack with associated tape drive provides 20M bytes of storage on a 1/4-inch cartridge tape for program load, patches, and translation. The tape drive provides an intelligent memory-mapped interface. It supports both streaming and edit modes and provides extensive error detection and correction capabilities, including the ability to correct very long burst errors.

3.3.3 Maintenance

The maintenance circuit pack uses microprocessors, VLSI, and digital signal processors to provide the following:

1. An RS-232-C interface to a hardwired maintenance/administration terminal (known as the system access terminal), and a low-level user interface in firmware that supplements the high-level interface in software.
2. A tip/ring interface to the remote maintenance center via the central office for automatic alarm reporting. It includes an autodialer, 212A modem emulation, and Level 2 X.25 protocol termination.
3. Cabinet environmental monitoring. If the temperature rises too high, the system is switched to power-fail transfer mode. When the temperature gradient across the cabinet increases to a predefined threshold, the user is reminded (via the system access terminal) to clean the air filters.
4. Power supply and battery holdover monitoring and control. The battery charger and power supplies are constantly monitored and controlled. On ac power failure, the entire system is powered from the batteries for ten seconds. Then the port carrier supplies are shut down and the control carrier is held over an additional ten minutes. Thus,

most commercial power outages are bridged without any interruption in service.

5. Power fail transfer control. As explained above, after ten seconds of battery holdover, the port carriers are shut down. At this time, selected voice terminals are connected directly (via relays) to central office trunks to provide emergency phone service.

The maintenance architecture is described in more detail in Ref. 13.

3.4 Extensions

The control complex can be extended by adding an interface to a multimaster System Bus (SBus), as shown in Fig. 8. In System 75, the SBus supports an additional processor with I/O that terminates the

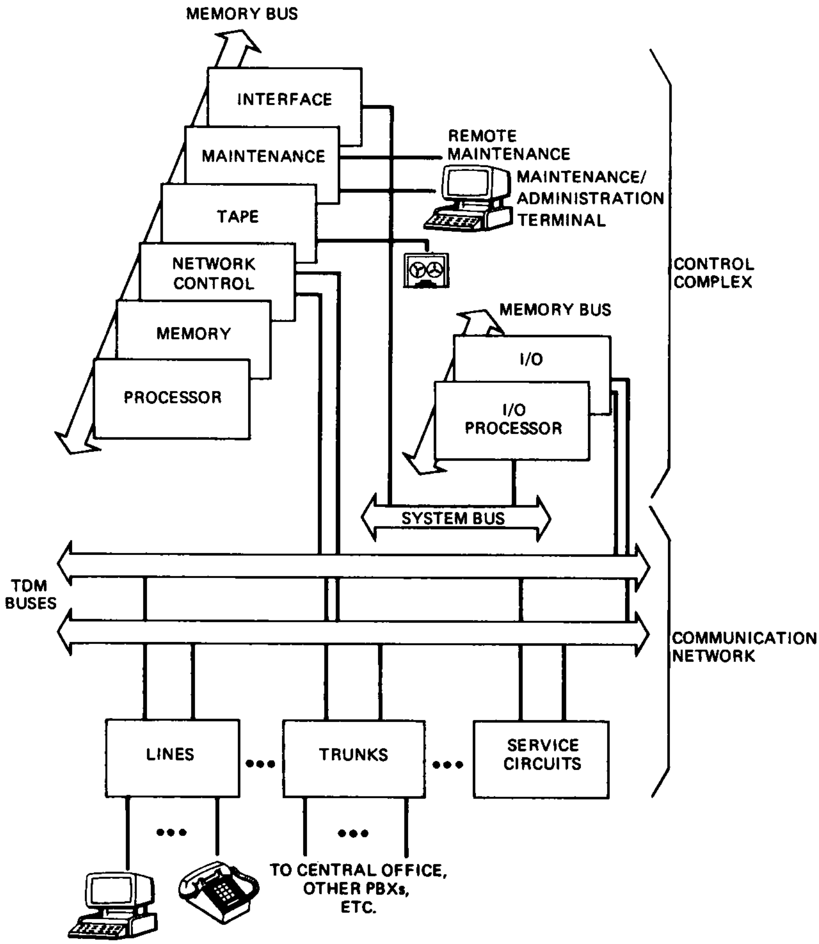

Fig. 8—System 75 communication and control architecture with I/O processor.

switched X.25 channels. These channels connect to adjunct systems such as the applications processor, or other nodes in a Distributed Communications Service (DCS) network.[14]

The I/O processor has its own 16-bit microprocessor plus 128K bytes of RAM and an SBus interface. It connects to the I/O interface through another short MBus. The I/O interface provides additional flexibility because it connects to the TDM buses and terminates the X.25 protocol and the underlying DCP protocol on all four channels. This permits the use of standard data switching, cabling, and termination features for a wide variety of system arrangements.

IV. SUMMARY

System 75 provides a digital communication network that serves the voice and data communication needs of medium-sized customers. Conferencing and a flexible gain plan are integral parts of the network. The control complex efficiently supports a modern operating-system-based software package, and can be expanded to support additional optional features. A wide range of station equipment is supported so that the system can be configured to fill the customer's needs in a cost-effective manner.

To protect the customer's investment, System 75 uses a flexible and highly modular architecture so that the system may be expanded to meet future needs. In addition, the modularity allows economical upgrading of the system as technology progresses.

REFERENCES

1. Bell Telephone Laboratories, *Transmission Systems for Communications*, Fifth edition, 1982, Chapter 28.
2. G. M. Anderson, J. F. Day, and L. A. Spindel, "A Communications Protocol for Integrated Digital Voice and Data Services in the Business Office," Proc. Sixth Int. Conf. on Computer Communication, London, September 1982, pp. 367–72.
3. H. G. Alles, "The Intelligent Communications Switching Network," IEEE Trans. Comm., *COM-27* (July 1979), pp. 1080–7.
4. L. A. Baxter, P. R. Berkowitz, and C. A. Buzzard, "Distributed Digital Conferencing System," U.S. Patent No. 4,389,720, June 21, 1983.
5. B. S. Moffitt and A. R. Ross, "Digital Conference Time Slot Interchanger," U.S. Patent No. 4,382,295, May 3, 1983.
6. B. S. Moffitt and A. R. Ross, "Multiport Memory Array," U.S. Patent No. 4,395,765, July 26, 1983.
7. "Digital Channel Bank Requirements and Objectives," AT&T Technical Reference PUB43801, December 1978.
8. "Integrated Services Digital Networks," Special Issue of IEEE Comm. Magazine, *22*, No. 1 (January 1984).
9. A. S. Loverde et al., "System 75: Physical Architecture and Design," AT&T Tech. J., this issue.
10. A. Meijer and P. Peeters, *Computer Network Architectures*, Rockville, MD: Computer Science Press, 1982, Chapter 4.
11. Special issue on "Digital Signal Processor," B.S.T.J., *60*, No. 7, Part 2 (September 1981).
12. B. S. Atal and J. R. Remde, "A New Model of Pulsed LPC Excitation for Producing Natural Sounding Speech at Low Bit Rates," Proc. ICASSP '82, Paris, France, May 1982, pp. 614–17.

13. K. S. Lu, J. D. Price, and T. L. Smith, "System 75: Maintenance Architecture," AT&T Tech. J., this issue.
14. R. S. Divakaruni, G. E. Saltus, and B. R. Savage, "New Directions in Enhanced Voice Networking," Proc. Sixth Int. Conf. on Computer Communication, London, September 1982, pp. 362–6.

AUTHORS

L. A. Baxter, B.S.E.E., 1975, Rochester Institute of Technology; M.S.E.E., 1977, University of Delaware; Bell Laboratories, 1977–1982; AT&T Information Systems Laboratories, 1983—. Mr. Baxter was initially involved with the exploratory design of office communication systems. Since 1980 he has worked on the design of the System 75 communications network. He currently is Supervisor of the System 75 Digital Switching Hardware group. Member, IEEE, Tau Beta Pi, Phi Kappa Phi.

P. R. Berkowitz, B.S.E.E., 1974, Columbia University; M.S.E.E., 1975, Columbia University; Bell Laboratories, 1975–1982; AT&T Information Systems Laboratories, 1983—. Mr. Berkowitz is Supervisor of the System 75 Circuit Design group. He was previously a member of the design team for the *Horizon*® communications system.

C. Alan Buzzard, B.S.E.E., 1964, M.S.E.E., 1965, Cornell University; Bell Laboratories, 1964–1982; AT&T Information Systems Laboratories, 1982–1983. Present affiliation Bell Communications Research, Inc. Mr. Buzzard's past responsibilities have included development of modems, data networks, and voice/data PBXs. His present interests are in speech coding, speech synthesis, and automatic speech recognition. Member, IEEE.

J. J. Horenkamp, B.S.E.E., 1964, St. Louis University; M.S.E.E., 1966, Columbia University; Doctor of Engineering Science, 1973, Columbia University; Bell Laboratories, 1964–1982; AT&T Information Systems Laboratories, 1983—. Mr. Horenkamp is Head of the System Design department responsible for hardware and firmware development and maintenance planning of System 75. He has been involved with exploratory development and final design of a variety of PBXs and customer premises telecommunication systems.

Frank E. Wyatt, B.S.E.E., 1969, M.S.E.E., 1970, University of Illinois; Bell Laboratories, 1971–1982; AT&T Information Systems Laboratories, 1983—. Mr. Wyatt has worked on a variety of business communication and management information systems. Since 1980 he has supervised the group that developed the control complex for System 75. Member, IEEE.