Chapter 28

Fault-Tolerant Design of Local ESS Processors[1]

W. N. Toy

Overview The stored program control of Bell System Electronic Switching Systems (ESS) has been under development since 1953. During this period, the No. 1 ESS, the No. 2 ESS, and the No. 3 ESS have been developed and used extensively by Bell System operating companies to provide commercial telephone service. These systems serve all types of telephone offices: The large-capacity No. 1 ESS serves metropolitan offices, the medium-capacity No. 2 ESS was designed for suburban offices, and the No. 3 ESS can be found in many small rural offices. The fault-tolerant design of ESS processors provides the same highly dependable telephone service established by the previous electromechanical systems. Pertinent process architecture features used to achieve ESS reliability objectives are discussed.

Introduction

Next to computer systems used in space-borne vehicles and U.S. defense installations, no other application has a higher availability requirement than a Bell System Electronic Switching System (ESS). These systems have been designed to be out of service no more than few minutes per year. Furthermore, design objectives permit no more than 0.01 percent of the telephone calls to be processed incorrectly [Downing, Nowak, and Tuomenoksa, 1964]. For example, when a fault occurs in a system, few calls in progress may be handled incorrectly during the recovery process.

At the core of every ESS is a single high-speed central processor [Harr, Taylor, and Ulrich, 1969; Browne et al., 1969; Staehler, 1977]. To establish an ultrareliable switching environment, redundancy of system components and duplication of the processor itself has been the approach taken to compensate for potential machine faults. Without this redundancy, a single component failure in the processor might cause a complete failure of the entire system. With duplication, a standby processor takes over control and provides continuous telephone service.

When the system fails, the fault must be quickly detected and isolated. Meanwhile, a rapid recovery of the call processing functions (by the redundant component(s) and/or processor) is necessary to maintain the system's high availability. Next, the fault must be diagnosed and the defective unit repaired or replaced. The failure rate and repair time must be such that the probability is very small for a failure to occur in the duplicated unit before the first one is repaired.

[1]*Subsetted from Proc. IEEE*, vol. 66, no. 10, October 1978, pp. 1,126–1,145.

Allocation and Causes of System Downtime

The outage of a telephone (switching) office can be caused by facilities other than the processor. While a hardware fault in one of the peripheral units generally results in only a partial loss of service, it *is* possible for a fault in this area to bring the system down. By design, the processor has been allocated two-thirds of the system downtime. The other one-third is allocated to the remaining equipment in the system.

Field experience indicates that system outages due to the processor may be assigned to one of four categories shown in Fig. 1 [Staehler and Watters, 1976]. The percentages in this figure represent the fraction of total downtime attributable to each cause. The four categories are as follows.

Hardware Reliability

Before the accumulation of large amounts of field data, total system downtime was usually assigned to hardware. We now know that the situation is more complex. Processor hardware actually accounts for only 20 percent of the downtime. With growing use of stored program control, it has become increasingly important to make such systems more reliable. Redundancy is designed into all subsystems so that the system can go down *only* when hardware failures occur simultaneously in duplicated units. However, the data now show that good diagnostic and trouble location programs are very critical parts of the total system reliability performance.

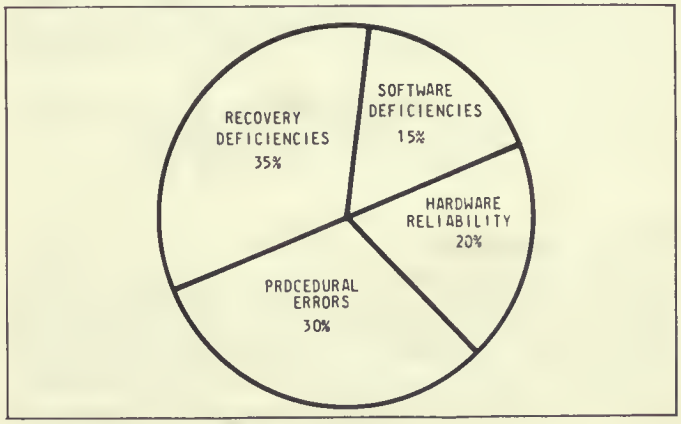

Fig. 1. System outage allocation.

Software Deficiencies

Software deficiencies include all software errors that cause memory mutilation, and program loops that can only be cleared by major reinitialization. Software faults are the result of improper translation or implementation of the original algorithm. In some cases, the original algorithm may have been incorrectly specified. Program changes and feature additions are continuously incorporated into working offices. Software accounts for 15 percent of the downtime.

Recovery Deficiencies

Recovery is the system's most complex and difficult function. Deficiencies may include the shortcomings of either hardware or software design to detect faults when they occur. When faults go undetected, the system remains extensively impaired until the trouble is recognized. Another kind of recovery problem can occur if the system is unable to properly isolate a faulty subsystem and configure a working system around it.

The many possible system states which may arise under trouble conditions make recovery a complicated process. Besides those already mentioned, unforeseen difficulties may be encountered in the field, and lead to inadequate recovery. Because of the large number of variables involved and because the recovery function is so strongly related to all other components of maintenance, recovery deficiencies account for 35 percent of the downtime.

Procedural Errors

Human error on the part of maintenance personnel or office administrators can also cause the system to go down. For example, someone in maintenance may mistakenly pull a circuit pack from the on-line processor while repairing a defective standby processor. Inadequate and incorrect documentation (e.g., users' manuals) may also be classified as human error. Obviously, the number of manual operations must be reduced if procedural errors are to be minimized. Procedural errors account for about 30 percent of the downtime.

The shortcomings and deficiencies of current systems are being continually corrected to improve system reliability.

Duplex Architecture

When a fault occurs in a nonredundant single processor, the system will remain down until the processor is repaired. In order to meet the ESS reliability requirement, *redundancy* is included in the system design; continuous and correct operation is maintained by duplicating all functional units within the processor. If one of the units fails, the duplicated unit is switched in, maintaining continuous operation. Meanwhile, the defective unit is repaired. Should a fault occur in the duplicated unit during the repair interval, the system will, of course, go down. If the repair interval is relatively short, the probability of simultaneous faults occurring in two identical units is quite small. This technique of redundancy has been used throughout each ESS.

The first-generation ESS processor structure consists of two store communities: program store (PS) and call store (CS). The program store is a read-only memory (ROM) containing the call processing, maintenance, and administration programs; it also contains long-term translation and system parameters. The call store contains the transient data related to telephone calls in progress. The memory is electrically alterable to allow its data to be changed frequently. In one particular arrangement, shown in Fig. 2, the complete processor is treated as a single functional block and is duplicated. This type of single-unit duplex system has two possible configurations: Either Processor 0 or Processor 1 can

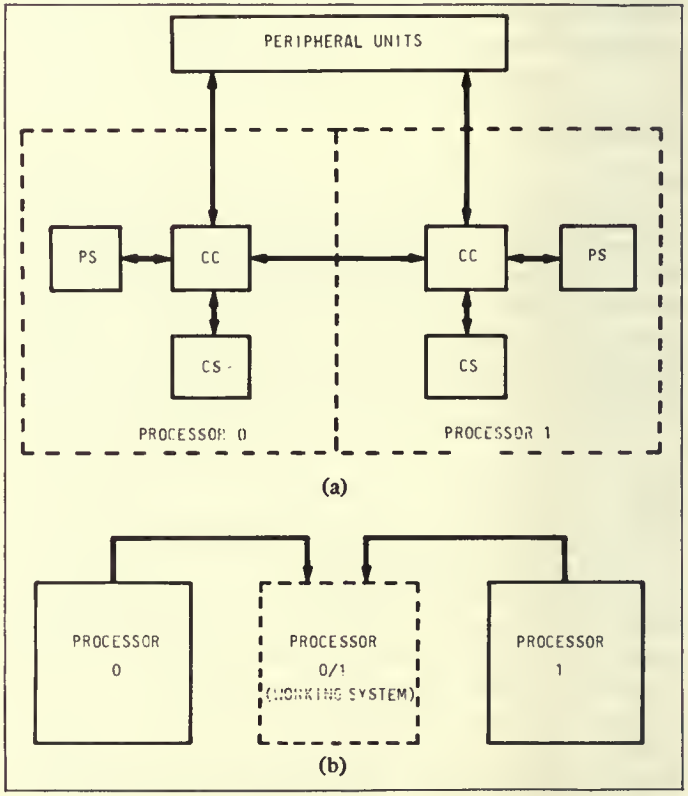

Fig. 2. Single-unit duplex configuration. *(a)* Processor structure. *(b)* Two possible configurations.

be assigned as the on-line working system, while the other unit serves as standby backup. The mean-time-to-failure (MTTF), a measure of reliability, is given by the following expression [Smith, 1972]:

$$\text{MTTF} = \frac{\mu}{2\lambda^2}$$

where μ is the repair rate (reciprocal of the repair time), and λ is the failure rate.

The failure rate (λ) of one unit is the summation of failure rates of all components within the unit. For medium and small ESS processors, Fig. 2 shows a system structure containing several functional units which are treated as a single entity, with λ still sufficiently small to meet the reliability requirement. The single-unit duplex configuration has the merit of being very simple in terms of the number of switching blocks in the system. This configuration simplifies not only the recovery program but also the hardware interconnection. It does this by eliminating the additional access required to make each duplicated block capable of switching independently into the on-line system configuration.

In the large No. I ESS, which contains many components, the MTTF becomes too low to meet the reliability requirement. In order to increase the value of the MTTF, either the number of components (failure rate) or the repair time must be reduced. Alternatively, the single-unit duplex configuration can be partitioned into a multiunit duplex configuration as shown in Fig. 3. In this arrangement, each subunit contains a smaller number of components and is able to be switched into a working system. The system will fail only if a fault occurs in the redundant subunit while the original is being repaired. Since each subunit contains fewer components, the probability of two simultaneous faults occurring in a duplicated pair of subunits is reduced. The MTTF of the multiunit duplex configuration can be computed by taking into consideration the conditional probability of a subunit failing during the repair time of the original subunit.

An example of a multiunit duplex configuration is shown in Fig. 3. A working system is configured with a fault-free CCx-CSx-CSBx-PSx-PSBx-PUBx arrangement, where x is either Subunit 0 or Subunit 1. This means there are 2^6, or 64 possible combinations of system configurations. The MTTF is given by the following expression:

$$\text{MTTF} = \frac{r\mu}{2\lambda^2}$$

where $r =$

$$\frac{1}{(\lambda_{CC}/\lambda)^2 + (\lambda_{CS}/\lambda)^2 + (\lambda_{CSB}/\lambda)^2 + (\lambda_{PS}/\lambda)^2 + (\lambda_{PSB}/\lambda)^2 + (\lambda_{PUB}/\lambda)^2}$$

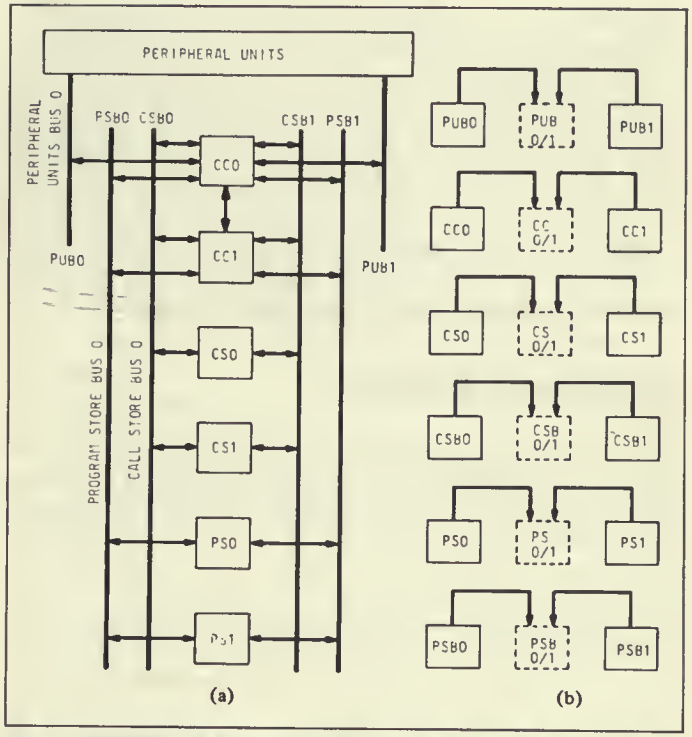

Fig. 3. Multiunit duplex configuration. *(a)* Processor structure. *(b)* Sixty-four possible configurations.

The factor r is at a maximum when the failure rate (λ_i) for each subunit is the same. In this case

$$\lambda_{CC} = \lambda_{CS} = \lambda_{CSB} = \lambda_{PS} = \lambda_{PSB} = \lambda_{PUB} = \lambda_I$$

or

$$\lambda_I = \frac{\lambda}{s}$$

where

s = number of subunits, $s = 6$, and $r = s$

At best, the MTTF is improved by a factor corresponding to the number of partitioned subunits. This improvement is not fully realized since equipment must be added to provide additional access and to select subunits. The partitioning of the subsystem into subunits as shown in Fig. 3 results in subunits of different sizes. Again, the failure rate for each individual subunit will not be the same; hence, the r-factor will be smaller than 6. Because of the relatively large number of components used in implementing

the No. 1 ESS, the system is arranged in the multiunit duplex configuration in order to meet the reliability requirement.

Reliability calculation is a process of predicting, from available failure rate data, the achieveable reliability of a system and the probability of meeting the reliability objectives for ESS applications. These calculations are most useful and beneficial during the early stages of design in order to assess various types of redundancy and determine the system's organization. In the small and medium ESS's, the calculations have supported the use of single-unit duplex structures. For large ESS's, it was necessary to partition the system into a multiunit duplex configuration.

Fault Simulation Techniques

One of the more difficult tasks of maintenance design is fault diagnosis. Its effectiveness in diagnostic resolution can be determined by simulation of the system's behavior in the presence of a specific fault. By means of simulation, design deficiencies can be identified and corrected prior to any system being deployed in the field. It is necessary to evaluate the system's ability to detect faults, to recover automatically back into a working system, and to provide diagnostic information where the fault is within a few replaceable circuit packs. Fault simulation, therefore, is an important aspect of maintenance design.

There are essentially two techniques used for simulating faults of digital systems: physical simulation or digital simulation. Physical simulation is a process of inserting faults into a physical working model. This method produces more realistic behavior under fault conditions. A wider class of faults can be applied to the system, such as a blown fuse or shorted backplane interconnection. However, fault simulation cannot begin until the design has been completed and the equipment is fully operational. Also, it is not possible to insert faults interior to an integrated circuit.

Digital fault simulation is a means of predicting the behavior under failure of a processor modeled in a computer program. The computer used to execute the program (the host) is generally different from the processor being simulated (the object). Digital fault simulation gives a high degree of automation and excellent access to interior points of logic to monitor the signal flow. It allows diagnostic test development and evaluation to proceed well in advance of unit fabrication. The cost of computer simulation can be quite high for a large, complex system.

The physical fault simulation method was first employed to generate diagnostic data for the Morris Electronic Switching System [Tsiang and Ulrich, 1962]. Over 50 000 known faults were purposely introduced into the central control to be diagnosed by its diagnostic program. Test results associated with each fault were recorded. They were then sorted and printed in dictionary format to formulate a trouble locating manual (TLM). Under trouble conditions, by consulting the TLM, it was possible to determine a set of several suspected circuit packs which might contain the defective component. Using the dictionary technique at the Morris system, the average repair time was kept low and maintenance was made much easier.

The experience gained in the physical fault simulation was applied and extended in the No. 1 ESS development [Downing, Nowak, and Tuomenoksa, 1964]. Each plug-in circuit pack was replaced by a fault simulator which introduced every possible type of single fault on the replaced package one at a time and then recorded the system reaction on magnetic tape. This was done for all circuit packs in the system. In addition to diagnostic data for dictionaries, additional data were collected to determine the adequacy of hardware and software in fault detection and system recovery. Deficiencies were corrected to improve the overall maintenance of the system.

A digital logic simulator called LAMP [Chang, Smith, and Walford, 1974] was developed for the No. 1A ESS development. It played an important role in the hardware and diagnostic development of the No. 1A Processor. The simulator is capable of simulating subsystem with as many as 65 000 logic gates. All classical faults for standard logic gates are simulatable with logic nodes stuck at "0" or stuck at "1." Before physical units are available, digital simulation can be very effective in verifying the design, evaluating diagnostic access, and developing diagnostic tests. Physical fault simulation has been demonstrated in the No. 1 ESS to give a very realistic behavior under fault conditions. The integration of both techniques was employed in the development of the No. 1A Processor to take advantages of both processes. The use of complementary simulation allows faults to be simulated physically (in the system laboratory) and logically (on a computer). Most of the deficiencies of one simulation process are compensated for by the other. The complementary method provided both a convenient method for validating the results and more extensive fault simulation data than would have been normally if either process were used individually. Fig. 4 shows the complementary process of fault simulation used in the No. 1A Processor development [Bowman et al., 1977; Goetz, 1974]. Maximum diagnostic performance was achieved from an integrated use of both simulation methods.

First Generation ESS Processors

The world's first ESS provided commerical telephone service at Morris, IL, in 1959 for about a year on a field trial basis [Keister, Ketchledge, and Lovell, 1960]. The system demonstrated the use of stored program control and the basic maintenance philosophy

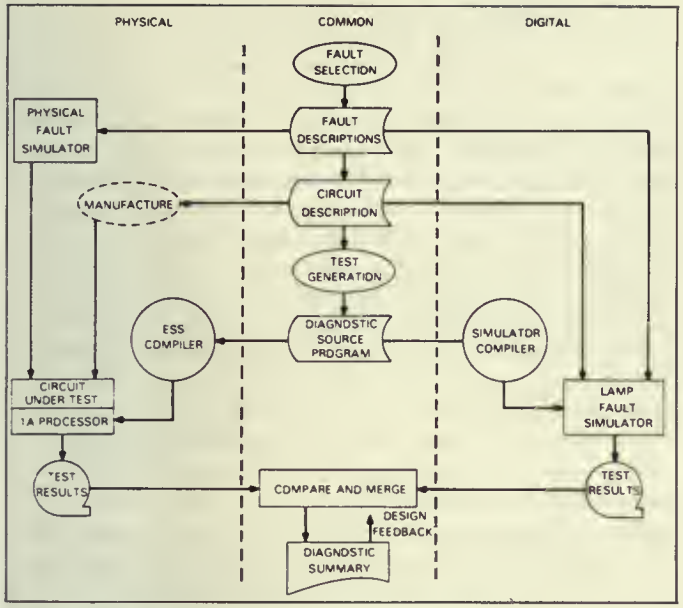

Fig. 4. Complementary fault-simulation system.

of providing continuous and reliable telephone service. The trial established valuable guides for designing a successor, the No. 1 ESS.

No. 1 ESS Processor

The No. 1 ESS was designed to serve large metropolitan telephone offices, ranging from several thousand to 65 000 lines [Keister, Ketchledge, and Vaughan, 1964]. As in most large switching systems, the processor represents only a small percentage of the total system cost. Therefore, performance and reliability were of primary importance in the design of the No. 1 Processor; cost was secondary. In order to meet the reliability standards established by electromechanical systems, all units essential to proper operation of the office are duplicated (see Fig. 3). The multiunit duplex configuration was necessary to increase the MTTF of the processor because of the large number of components in each of the functional blocks.

Even with duplication, troubles must be found and corrected quickly to minimize exposure to system failure due to multiple troubles. All units are monitored continually so that troubles in the standby units are found just as quickly as those in the on-line units. This is accomplished by running the on-line and standby units in the synchronous and match mode of operation [Downing, Nowak, and Tuomenoksa, 1964]. Synchronization requires that clock timing signals be in close tolerance so that every operation in both halves is performed in step, and key outputs are compared for error detection. The synchronization of duplicated units is accomplished by having the on-line oscillator output drive both clock circuits. There are two match circuits in each central control (CC). Each matcher compares 24 bits within one machine cycle of 5.5 μs. Fig. 5 shows that each matcher has access to six sets of internal nodes (24 bits per node). In the routine match mode, the points matched in each cycle are dependent upon the instruction being executed. The selected match points are those most pertinent to the data processing steps occurring during a given machine cycle. The two matchers in each CC compare the same sets of selected test points. The capability of each CC to compare a number of internal nodes provides a highly effective means of detecting hardware errors.

If a mismatch occurs, an interrupt is generated, which causes the fault-recognition program to run. The basic function of this program is to determine which half of the system is faulty. The

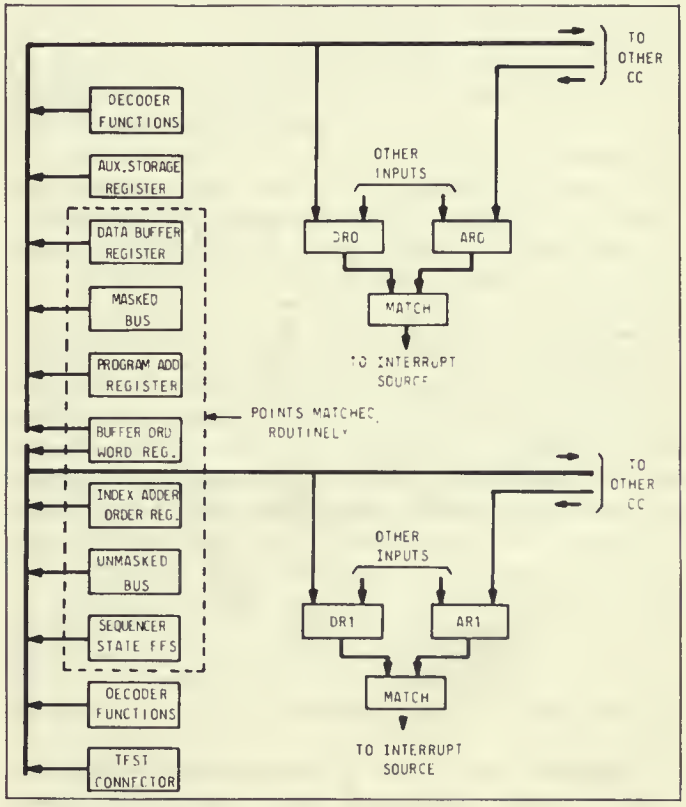

Fig. 5. No. 1 CC match access.

suspected unit is removed from service and the appropriate diagnostic program is run to pinpoint the defective circuit pack.

The No. 1 ESS was designed during the discrete component era (early 1960's) using individual components to implement logic gates [Cagle et al., 1964]. The CC contains approximately 12 000 logic gates. Although this number appears small when compared to large-scale integration (LSI) technology, the No. 1 Processor was a physically large machine for its time.

The match circuits capable of comparing internal nodes are the primary tools incorporated into the CC for diagnosing as well as detecting troubles. Specified information can be sampled by the matchers and retained in the match registers for examination. This mode of operation obtains critical data during the execution of diagnostic programs.

The early program store used permanent magnet twister (PMT) modules as basic storage elements [Ault et al., 1964]. They are a form of ROM in which system failures cannot alter the information content. Experience gained from the Morris field test system, which used the less reliable flying spot store, indicated that Hamming correction code was highly effective in providing continuous operation. At the time of development, it was felt that PMT modules might not be reliable enough. Consequently, the program store word included additional check bits for single-bit error correction (Hamming code). In addition, an overall parity check bit which covers both the data and their address is included in the word. The word size consists of 37 bits of information and seven check bits. When an error is corrected during normal operation, it is logged in an error counter. The maintenance program has access to this counter. Also, detection of a single error in the address or a double error in the word will cause an automatic retry.

The call store is the temporary read and write memory for storing transient data associated with call processing. Ferrite sheet memory modules are the basic storage elements used in implementing the call store in the No. 1 ESS [Genke, Harding, and Staehler, 1964]. The call store used in most No. 1 offices is smaller than the program store. (At the time of design, the cost per bit of call store was considerably higher than that of program store.) Also, ferrite sheet memory modules were considered to be very reliable devices. Consequently, single-bit error detection rather than Hamming correction code was provided in the call store.

There are two parity check bits: one over both the address and data, and the other over the address only. Again, as in the program store, automatic retry is performed whenever an error is detected, and the event is logged in an error counter for diagnostic use.

Troubles are normally detected by fault-detection circuits, and error-free system operation is recovered by fault recognition programs [Downing, Nowak, and Tuomenoksa, 1964]. This requires the on-line processor to be capable of making a proper decision. If this is not possible, an emergency action timer will "time out" and activate special circuits to establish various combinations of subsystems into a system configuration. A special program which is used to determine whether or not the assembled processor is sane takes the processor through a series of tests arranged in a maze. Only one correct path through the maze exists. If the processor passes through successfully, the timer will be reset, and recovery is successful. If recovery is unsuccessful, the timer will time out again, and the rearrangement of subsystems will be tried one at a time (e.g., combinations of CC, program store, and program store bus systems). For each selected combination, the special sanity program is started and the sanity timer is activated. This procedure is repeated until a working configuration is found. The sanity program and sanity timer determine if the on-line CC is functioning properly. The active CC includes the program store and the program store bus.

Operational Results of No. 1 ESS

The No. 1 ESS has been in commercial operation since 1965. Over 1000 systems are providing telephone service to more than 15 million subscribers. The performance of the No. 1 ESS has continually improved over a decade of continued effort to improve all phases of software and hardware.

Fig. 6 shows the result of field data accumulated over many machine operating hours. This curve was derived from data in a paper [Fleckenstein, 1974] presented at the 1974 International Switching Symposium in Munich, Germany, and data supplied by W. C. Jones of Bell Laboratories.

When the No. 1 ESS was first put into commercial service, many outages occurred because of software and hardware inadequacies that could only be weeded out with field experience. The

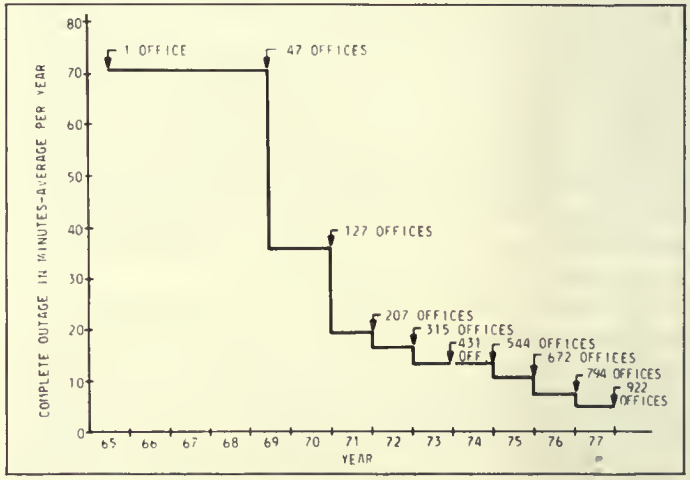

Fig. 6. No. 1 ESS service performance.

inexperience of maintenance personnel also contributed heavily towards system outages. Most hardware and software bugs were corrected during the early years of operation. However, deficiencies still exist, and designs are continually upgraded in working systems. Continual improvements include better diagnostic access, more complete fault recognition and isolation programs, and more effective system recovery.

Improved diagnostic capability reduces repair time and human errors by decreasing the amount of human interaction required by the machine. Better maintenance procedures and more experienced craftpersonnel also contribute to improved system performance. The curve in Fig. 6 shows that the outage rate improved as machine design and operating personnel matured.

No. 2 ESS Processor

The No. 2 ESS was developed during the mid-1960's [Spencer and Vigilante, 1969]. This system was designed for medium-size offices ranging from 1000 to 10 000 lines. The processor's design was derived from experience with the common stored program control of a private branch exchange (PBX), the No. 101 ESS [Seley and Vigilante, 1964]. Since the capacity requirement of the No. 2 ESS was to be less than that of the No. 1 ESS, cost became one of the more important design considerations. (Reliability is equally important in all systems.) The No. 2 ESS contains much less hardware than the No. 1 ESS. Understandably, its component failure rate is also substantially less. Its CC contains approximately 5000 gates (discrete components). To reduce cost and increase reliability, resistor-transistor logic (RTL) gates were chosen for the No. 2's processor since resistors are less expensive and more reliable than diodes [the No. 1 Processor used diode-transistor logic (DTL)].

Because the No 2's CC, program store, and call store are smaller, they are grouped together as a single switchable block in the single-unit duplex configuration shown in Fig. 2. Calculations indicate that its MTTF is approximately the same as the No. 1 multiunit duplex structure, with each of the functional blocks and associated store buses grouped together as a switchable block. The use of only two system configurations reduces considerably the amount of hardware needed to provide gating paths and control for each functional unit. Moreover, the recovery program is simplified, and the reliability of the system is improved.

The No. 2 Processor runs in the synchronous and match mode of operation [Beuscher et al., 1969]. The on-line oscillator output drives both clock circuits in order to keep the timing synchronized. The match operation is not as extensive as it is in the No. 1 ESS. For simplicity, there is only one matcher in the No. 2 ESS; it is located in the nonduplicated maintenance center (see Fig. 7). The matcher always compares the call store input registers in the two CC's when call store operations are performed synchronously. A fault in almost any part of either CC quickly results in a call store input register mismatch. This occurs because almost all data manipulation performed in both the program control and the input-output (I/O) control involves processed data returning to the call store. The call store input is the central point whereby data eventually funnel through to the call store. By matching the call store inputs, an effective check of the system equipment is provided. Compared to the more complex matching of the No. 1 Processor, error detection in the No. 2 Processor may not be as instantaneous since only one crucial node in the processor is matched. Certain faults in the No. 2 Processor will go undetected until the errors propagate into the call store. This interval is probably no more than tens or hundreds of microseconds. During such a short interval, the fault would affect only a single call.

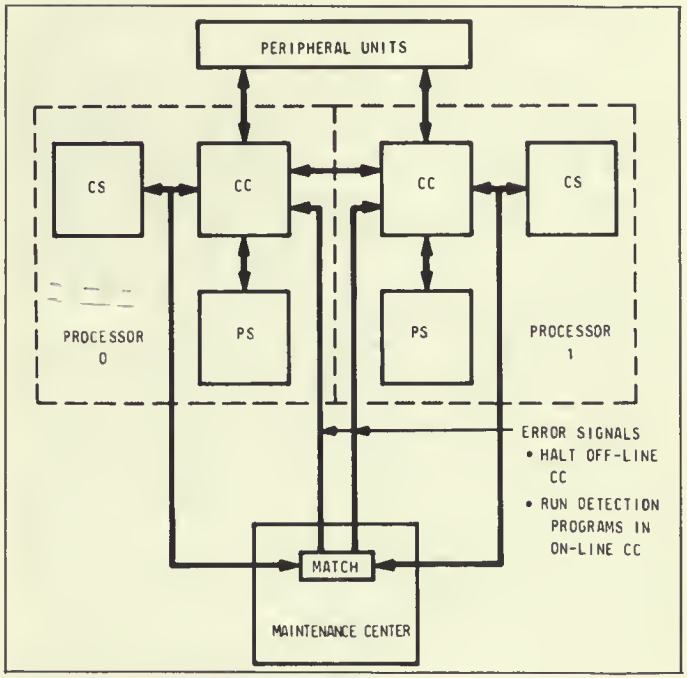

Fig. 7. No. 2 CC match access.

The No. 2 ESS matcher is not used as a diagnostic tool as is the matcher in the No. 1 Processor. Therefore, additional detection hardware is designed into the No. 2 Processor to help diagnose as well as detect faults.

When a mismatch occurs, the detection program is run in the on-line CC to determine if it contains the fault. This is done while the standby processor is disabled. If a solid fault in the on-line processor is detected by the mismatch detection program, the control is automatically passed to the standby processor, causing it to become the on-line processor. The faulty processor is disabled

and diagnostic tests are called in to pinpoint the defective circuit pack.

The program store also uses PMT modules as basic storage elements, with a word size of 22 bits, half the width of the No. 1's word size. Experience gained in the design and operation of the No. 101 ESS (PBX) showed that PMT stores were very reliable. The additional protection provided in the No. 1 Processor against memory faults by error correction was not considered to be as essential in the No. 2 Processor. This and the need to keep the cost down led to the choice of error detection *only* instead of the more sophisticated Hamming correction code.

Error detection works as follows: one of the 22 bits in a word is allocated as a parity check bit. The program store contains both program and translation data. Additional protection is provided by using odd parity for program words and even parity for translation data. This detects the possibility of accessing the translation data area of memory as instruction words. For example, a software error may cause the program to branch into the data section of the memory and execute the data words as instruction words. The parity check would detect this problem immediately. The program store includes checking circuits to detect multiple-word access. Under program control, the sense amplifier threshold voltage can be varied in two discrete amounts from its nominal value to obtain a measure of the operating margin. The use of parity check was the proper choice for the No. 2 ESS in view of the high reliability of these memory devices.

The No. 2 Processor call store uses the same ferrite sheet memory modules as the No. 1 Processor. However, the No. 2's data word is 16 bits wide instead of 24. Fault detection depends heavily upon the matching of the call store inputs when the duplex processors run in the synchronous mode. Within the call store circuit, the access circuitry is checked to see that access currents flow in the right direction at the correct time and that only two access switches are selected in any store operation. This ensures that only one word is accessed in the memory operation. Similarly, threshold voltages of the sense amplifiers may be varied under program control to evaluate the operating margins of the store. No parity check bit is provided in the call store.

Each processor contains a program timer which is designed to back up other detection methods. Normally, the on-line processor clears the timer in both processors at prescribed intervals if the basic call processing program cycles correctly. If, however, a hardware or software trouble condition exists (e.g., a program may go astray or a long program loop may prevent the timer from being cleared), the timer will time out and automatically produce a switch. The new on-line processor is automatically forced to run an initialization restart program which attempts to establish a working system. System recovery is simplified by using two possible system configurations rather than the multiunit duplex system.

Second Generation of ESS Processors

The advent of silicon integrated circuits (IC's) in the mid-1960's provided the technological climate for dramatic miniaturization, improved performance, and cost-reduced hardware. "1A technology" refers to the standard set of (IC) devices, apparatus, and design tools that were used to design the No. 1A Processor and the No. 3A Processor [Becker et al., 1977]. The choice of technology and the scale of integration level was dictated by the technological advances made between 1968 and 1970. Small-scale integration (SSI), made possible by bipolar technology, was capable of high yield production. Because of the processor cycle time, high-speed logic gates with propagation delays from 5 to 10 ns were designed and developed concurrent with the No. 1A Processor.

No. 1A Processor

The No. 1A Processor, successor to the No. 1 Processor, was designed primarily for the control of large local and toll ESS with high processing capabilities (the No. 1A ESS and No. 4 ESS, respectively) [Budlong et al., 1977]. An important objective in developing the No. 1A ESS was to maintain commonality with the No. 1 ESS. High capacity was achieved by implementing the new No. 1A integrated technology and a newly designed system structure. These changes made possible an instruction execution rate that is four to eight times faster than the No. 1 Processor. Compatability with the No. 1 ESS also allows the No. 1A Processor to be retrofitted into an in-service No. 1 ESS, replacing the No. 1 Processor when additional capacity is needed. The first 1A Processor was put into service in January 1976, as control for a No. 4 ESS in Chicago. Less than one year later, the first No. 1A ESS was put into commercial operation. By 1980, several hundred will be in service [Nowak, 1976].

The No. 1A Processor architecture is similar to its predecessor in that all of its subsystems have redundant units and are connected to the basic CC via redundant bus systems [Bowman et al., 1977]. One of the No. 1A Processor's major architectural differences is its program store [Ault et al., 1977]. It has a writable random-access memory (RAM) instead of PMT ROM. By combining disk memory and RAM, the system has the same amount of memory as a system with PMT, but at a lower cost. Backup copy of program and translation data is kept on disk. Other programs (e.g., diagnostics) are brought to RAM as needed; the same RAM spare is shared among different programs. More important is the system's ability to change the content of the store quickly and automatically. This simplifies considerably the administration and updating of program and translation information in working offices.

The additional disk (file store) subsystem adds flexibility to the

No. 1A Processor [Ault et al., 1977], but it also increases the complexity of system recovery. Fig. 8 shows the multiunit duplex 1A Processor. This configuration is similar to the No. 1 Processor arrangement (see Fig. 3) with a duplicated file store included. The file store communicates with the program store or call store via the CC and the auxiliary unit bus. This allows direct memory access between the file store and the program store or the call store. The disk file and the auxiliary unit bus are grouped together as a switchable entity.

Error detection is achieved by the duplicated and matched synchronous mode of operation, as in the No. 1 Processor. Both CC's operate in step and perform identical operations. The matching is done more extensively in the 1A to obtain as complete a check as possible. There are two match circuits in each processor. Each matcher has the ability to compare 24 internal bits to 24 bits in its mate once every machine cycle. (A machine cycle is 700 ns.) Any one of 16 different 24-bit internal nodes can be selected for comparison. The choice is determined by the type of instruction being executed. Rather than compare the same nodes in both CC's, the on-line and the standby CC's are arranged to match different sets of data. Four distinct internal groups are matched in the same machine cycle. This ensures the correct execution of any instruction.

The No. 1A Processor design is an improvement of the No. 1 Processor design. The No. 1A Processor incorporates much more checking hardware throughout various functional units in addition to matching hardware. Checking hardware speeds up fault detection and also aids the fault recovery process by providing indications that help isolate the faulty unit. The matching is used in various modes for maintenance purposes. This capability provides powerful diagnostic tools in isolating faults.

The program store and call store use the same hardware technology. The CC contains approximately 50 000 logic gates. While the initial design of the stores called for core memories, they have been replaced with semiconductor dynamic MOS memories. The word size is 26 bits: 24 data bits and two parity check bits. In the No. 1 Processor, the program store and the call store are fully duplicated. Because of their size, duplication requires a considerable amount of hardware, resulting in higher cost and increased component failures. To reduce the amount of hardware in the No. 1A Processor's store community, the memory is partitioned into blocks of 64K words, as shown in Fig. 9. Two additional store blocks are provided as roving spares. If one of the program stores fails, a roving program store spare is substituted and a copy of the program in the file store is transferred to the program store replacement. This type of redundancy has been made possible by the ability to regenerate data stored in a failing unit. Since a program store can be reloaded from the file store in less than a second, a roving spare redundancy plan is sufficient to meet the reliability requirement. As a result, Hamming correction code was not adopted in the No. 1A program store. However; it is essential that an error be detected quickly. Two parity check bits are generated over a partially overlapped, interleaved set of

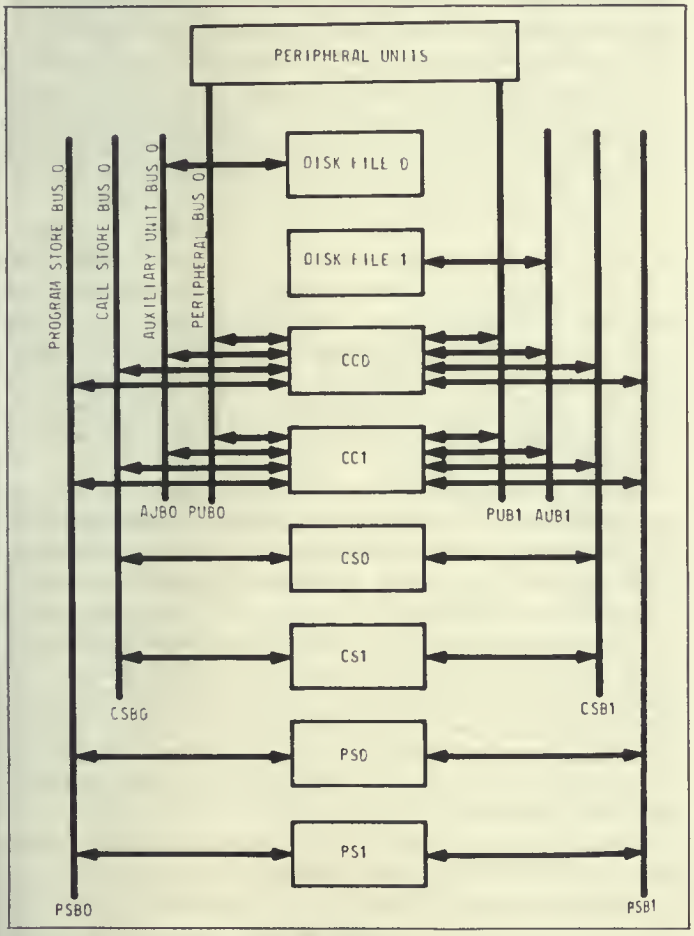

Fig. 8. No. 1A processor configuration.

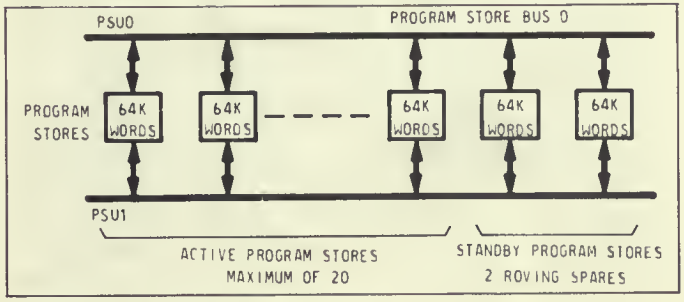

Fig. 9. No. 1A program store structure.

data bits and address. This overlapping is arranged to cope with particular memory circuit failures which may affect more than one bit of a word.

The 1A call stores contain both translation data backed up on the file stores and call-related transient data which are difficult to regenerate. The roving spare concept is expanded for the call stores to include sufficient spares to provide full duplication of transient data. If a fault occurs in a store that contains translation data, one of the duplicated stores containing transient call data is preempted and loaded with the necessary translation data from the duplicated copy in the file store. A parity check is done in the same manner as in the program store, using two check bits.

The combination of writable program store and file store provides a very effective and flexible system architecture for administrating and implementing a wide variety of features which are difficult to obtain in the No. 1 ESS. However, this architecture also complicates the process of fault recognition and recovery. Reconfiguration into a working system under trouble conditions is an extensive task, depending on the severity of the fault. (For example, it is possible for the processor to lose its sanity or ability to make proper decisions.) An autonomous hardware processor configuration (PC) circuit is provided in each CC to assist in assembling a working system. The PC circuit consists of various timers which ensure that the operational, fault recovery, and configuration programs are successfully executed. If these programs *are not* executed, the PC circuit controls the CC-to-program memory configuration, reloading program memory from file store when required, and isolating various subsystems from the CC until a working system is obtained.

No. 3A Processor

The No. 3A Processor was designed to control the small No. 3 ESS [Irland and Stagg, 1974], which can handle from 500 to 5000 lines. One of the major concerns in the design of this ESS was the cost of its processor. The low cost and high speed of integrated logic circuitry made it possible to design a cost-effective processor that performed better than its discrete component predecessor, the No. 2 Processor. The No. 3A project was started in early 1971. The first system cut into commercial service in late 1975.

Because the number of components in the No. 3A Processor is considerably less than in the No. 1A Processor, all subsystems are fully duplicated, including the main store. The CC, the store bus, and the store are treated as a single switchable entity rather than individual switchable units as in the No. 1A Processor. The system structure is similar to the No. 2 ESS. Experience gained in the design and operation of the No. 2 provided valuable input for the No. 3 Processor design.

The 3A's design makes one major departure from previous ESS processor designs: it operates in the nonmatched mode of duplex operation. The primary purpose of matching is to detect errors. A mismatch, however, does not indicate *where* (which one of the processors) the fault has occurred. A diagnostic fault-location program must be run to localize the trouble so that the defective unit can be taken off-line. For this reason, the No. 3A Processor was designed to be self-checking, with detection circuitry incorporated as an integral part of the processor. Faults occurring during normal operation are discovered quickly by detecting hardware. This eliminates the need to run the standby system in the synchronous and match mode of operation, or the need to run the fault recognition program to identify the defective unit when a mismatch occurs.

The synchronous and match mode arrangement of the No. 1 Processor and the No. 2 ESS provides excellent detection and coverage of faults. However, there are many instances (e.g., periodic diagnostics, adminstration changes, recent change updates, etc.) when the system is not run in the normal match mode. Consequently, during these periods, the system is vulnerable to faults which may go undetected. The rapid advances in integrated circuit technology make possible the implementation of self-checking circuits in a cost-effective manner. This eliminates the need for the synchronous and match mode of operation. Self-checking design is covered in more detail in Toy [1978].

Another new feature in ESS processor design is the application of microprogram technique in the No. 3A [Storey, 1976]. This technique provides a regular procedure of implementing the control logic. Standard error detection is made part of the hardware to achieve a high degree of checkability. Sequential logic, which is difficult to check, is easily implemented as a sequence of microprogram steps. Microprogramming offers many attractive features: it is simple, flexible, easy to maintain, and easy to expand.

The No. 3A Processor paralleled the design of the No. 1A Processor in its use of an electrically alterable (writable) memory. However, great strides in semiconductor memory technology after the No. 1A became operational permitted the use of semiconductor memory in the 3A rather than core memory.

The 3A's call store and program store are consolidated into a single store system. This reduces cost by eliminating buses, drivers, registers, and controls. A single store system no longer allows concurrent access of call store and program store. However, this disadvantage is more than compensated for by the much faster semiconductor memory. Its access time is 1 μs (the earlier PMT stores had an access time of 6 μs).

Normal operation requires the on-line processor to run and process calls while the standby processor is in the halt state, with its memory updated for each write operation. For the read operation, only the on-line memory is read, *except* when a parity error occurs during a memory read. This results in a micropro-

gram interrupt, which reads the word from the standby store in an attempt to bypass the error.

As discussed previously, the No. 2 Processor (first generation) is used in the No. 2 ESS for medium-size offices. It covers approximately 4000 to 12 000 lines, with a call handling capability of 19 000 busy-hour calls. (The number of calls is related to the calling rate of lines during the busy hour.) The microprogram technique used in the No. 3A Processor design allows the No. 2 Processor's instruction set to be emulated. This enables programs written in the No. 2 assembly language to be directly portable to the No. 3A Processor. The ability to preserve the call processing programs permits the No. 2 ESS to be updated with the No. 3A Processor without having to undergo a complete, new program development.

The combination of the No. 3A Processor and the peripheral equipment of the No. 2 ESS is designated as the No. 2B ESS. It is capable of handling 38 000 busy-hour calls, twice the capability of the No. 2 ESS [Mandigo, 1976]. The No. 2B ESS can be expanded to cover about 20 000 lines. Furthermore, when an existing No. 2 ESS system in the field exceeds its real-time capacity, the No. 2 Processor can be taken out and replaced with the No. 3A Processor. The retrofit operation has been carried out successfully in working offices without disturbing telephone service.

Summary

In order to achieve the reliability requirements, all ESS subsystem units are duplicated. When a hardware failure occurs in any of the subunits, the processor is reconfigured into a working system around the defective unit. The partitioning of subsystem units into switching blocks varies with the size of the ESS processors. For the medium- or small-size processors such as the No. 2 or the No. 3, the central control, the main memory, the bulk memory, and the store bus are grouped as a single switchable entity. A failure in one of the subunits is considered a failure in the switchable block. Since the number of components within a switchable block is sufficiently small, this type of single-unit duplex configuration meets the reliability requirement. For larger processors such as the No. 1 or the No. 1A, the central control, the program store, the call store, the store buses, and the bulk file store are treated individually as switchable blocks. This multiunit duplex configuration allows a considerable number of combinations in which a working system can be assembled. The system is down only when two simultaneous failures occur, one in the subunit and the other in the duplicated subunit. A greater fault tolerance is possible with this configuration. This type of configuration is necessary for the large processor because each subunit contains a larger number of components.

The first generation of ESS processors, which includes the No. 1 and the No. 2, have provided commerical service since 1965 and 1969, respectively. The No. 1 ESS serves large telephone offices (metropolitan); the No. 2 is used in medium-size offices (suburban). Their reliability requirements are the same. Both processors depend on integrated maintenance software, with the hardware that must (1) quickly detect a system failure condition, (2) isolate and configure a working system around the faulty subunit, (3) diagnose the faulty unit, and (4) assist the maintenance personnel in repairing the unit. The primary detection technique is the synchronous and match mode of operation of both central controls. Matching is done more extensively in the No. 1 than in the No. 2 since cost is one of major considerations in the design of the No. 2 Processor. In addition to matching, coding techniques, diagnostic access, and other check logic have been incorporated into the basic design of these processors to realize the reliability objectives.

The widespread acceptance of the No. 1 ESS and the No. 2 ESS has created the need for a second generation of ESS processors: the No. 1A and the No. 3A. They offer greater capability and are also more cost-effective. Both processors use the same integrated technology. The 1A Processor extends its performance range by a factor of four to eight times over the No. 1 Processor by using faster logic and faster memory. The 1A design takes advantage of the experience gained in the design and operation of the No. 1 ESS. The No. 1A Processor provides considerably more hardware for error detection and more extensive matching of a large number of internal nodes within the central control. The design of the No. 3A Processor had benefited by the experience gained from the No. 2 ESS. A major departure in the design of the 3A Processor from the design of other ESS processors is the nonsynchronous and the nonmatch mode of operation. The No. 3A Processor uses self-checking as primary means of error detection. Another departure is in the design of the No. 3A Processor's control section; it is microprogrammed. The No. 3A Processor's flexibility permits emulation of the No. 2 Processor quite easily.

References

Ault et al. [1977]; Ault, Gallaher, Greenwood, and Koehler [1964]; Becker et al. [1977]; Beuscher et al. [1969]; Bowman et al. [1977]; Browne, Quinn, Toy, and Yates [1969]; Budlong et al. [1977]; Cagle et al. [1964]; Chang, Smith, and Walford [1974]; Downing, Nowak, and Tuomenoksa [1964]; Fleckenstein [1974]; Genke, Harding, and Staehler [1964]; Goetz [1974]; Harr, Taylor, and Ulrich [1969]; Irland and Stagg [1974]; Keister, Ketchledge, and Lovell [1960]; Keister, Ketchledge, and Vaughan [1964]; Mandigo [1976]; Nowak [1976]; Seley and Vigilante [1964]; Smith [1972]; Spencer and Vigilante [1969]; Staehler [1977]; Staehler and Watters [1976]; Storey [1976]; Toy [1978]; Tsiang and Ulrich [1962].