No. 1 ESS ADF:

Message Store—A Disk Memory System

By E. J. AITCHESON, C. F. AULT and R. G. SPENCER

(Manuscript received June 2, 1970)

A high thruput, 60 megabit, fixed head per track disk memory system was designed to meet the bulk message store requirements for the No. 1 ESS ADF System. The application in this store and forward message switcher demanded that the disk store meet the stringent reliability requirements consistent with those of an electronic switching system. The article describes the system design features, a temperature-compensated read/write scheme, a synchronizing motor servo system, and automatic maintenance techniques.

I. INTRODUCTION

An essential element of the No. 1 ESS ADF store and forward data handling system is an economical, nondestructive readout, bulk memory with an average thruput of 153 kb/s. Core and ferrite sheet memories meet the access requirement, but their costs and destructive readout make them less attractive than the slower access disk or drum large capacity memories.[1]

In a study conducted to select a memory file that would best fit the ADF application, the following requirements were established for the memory:

(*i*) present commercial production and indication of continuing supply over the years,
(*ii*) performance over the expected 40-year life of the system with minimum downtime and repair cost,
(*iii*) fixed head/track access with a minimum of mechanical equipment, and
(*iv*) capacity of at least 50 megabits.

The file selected, a Burroughs Corporation model BC475 Disk File, has a fixed head/track access scheme, a capacity of 60 megabits, and an average access time of 20 milliseconds (1/2 of a revolution).

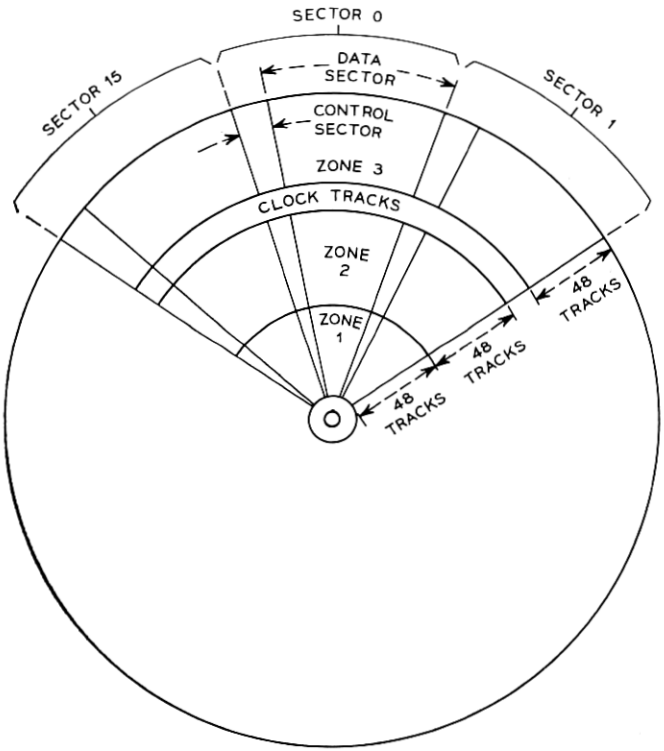

Fig. 1—Disk surface geometry as determined by word clock.

The average message thruput requirement for the No. 1 ESS ADF application was obtained by partitioning the disk faces into 16 pie-shaped sections (Fig. 1) and by the application of queueing techniques. Thus, during a given disk revolution, a maximum of 16 blocks by 32 words/block by 24 bits/word can be transferred to or from any given disk address providing a maximum thruput of 308 kb/s.

The disk file is incorporated in the message store subsystem, which consists of a logic controller, access circuits, synchronizing servo system, and the 60-megabit disk file. The logic controller, access circuits, and servo were designed at Bell Telephone Laboratories and manufactured by the Western Electric Company. Important design parameters are tabulated in Table I. The controller was designed with No. 1 ESS technology and used available circuit packs, wherever possible; thus, the unit meets No. 1 ESS circuit reliability objectives.[2] The

design also includes self-checking maintenance hardware and sufficient access to allow program-controlled maintenance testing consistent with the ESS philosophy.[3] Maintenance programs executed from the common-control ADF central processor detect and isolate faulty components to within three replaceable, plug-in circuit packs (Fig. 2).

Desired reliability and message protection are achieved by fully duplicating the message store subsystem and running the two halves in a synchronous match mode. Duplicate communication buses are provided so that the units can be kept in step and configured in a few microseconds, by program control, in the presence of a fault. This provides maximum protection against mutilation or loss of call data.

II. DISK FILE

The file, motor drive, and head pressure system are manufactured by the Burroughs Corporation. Each file contains four 26-inch diameter brass disks (eight memory faces), plated with a thin film of nickel cobalt. The disks are belt driven on a common shaft from a central pedestal position. To prevent contamination, the disks and read/write heads are sealed in a dust-tight enclosure maintained at a pressure of one inch of water by an external blower-filter arrangement. The disk file incorporates a fixed head/track access matrix with the heads held in near proximity to the surface by a pneumatically driven head piston working against the counter force of a laminar air cushion and retracting spring.

The file and dust-sealed enclosure are mounted in an environmentally controlled temperature chamber held at an elevated temperature with respect to the office ambient. This provides file stability over the office operating range of 32° to 120°F.

TABLE I—MESSAGE STORE CHARACTERISTICS

Controller power	840W @ +24VDC, 960W @ −48VDC
Motor inverter power input	1300 W
Motor speed	1725 rpm
Motor type	2 hp, 208 volts, 3 ϕ (inverter powered)
Data packing density	1100 bits/inch
Type of recording	NRZ
Bits/file	60 megabits
Writing surfaces	8 (4 disks) 1/8″ thick 26″ diameter
Number data tracks	1152

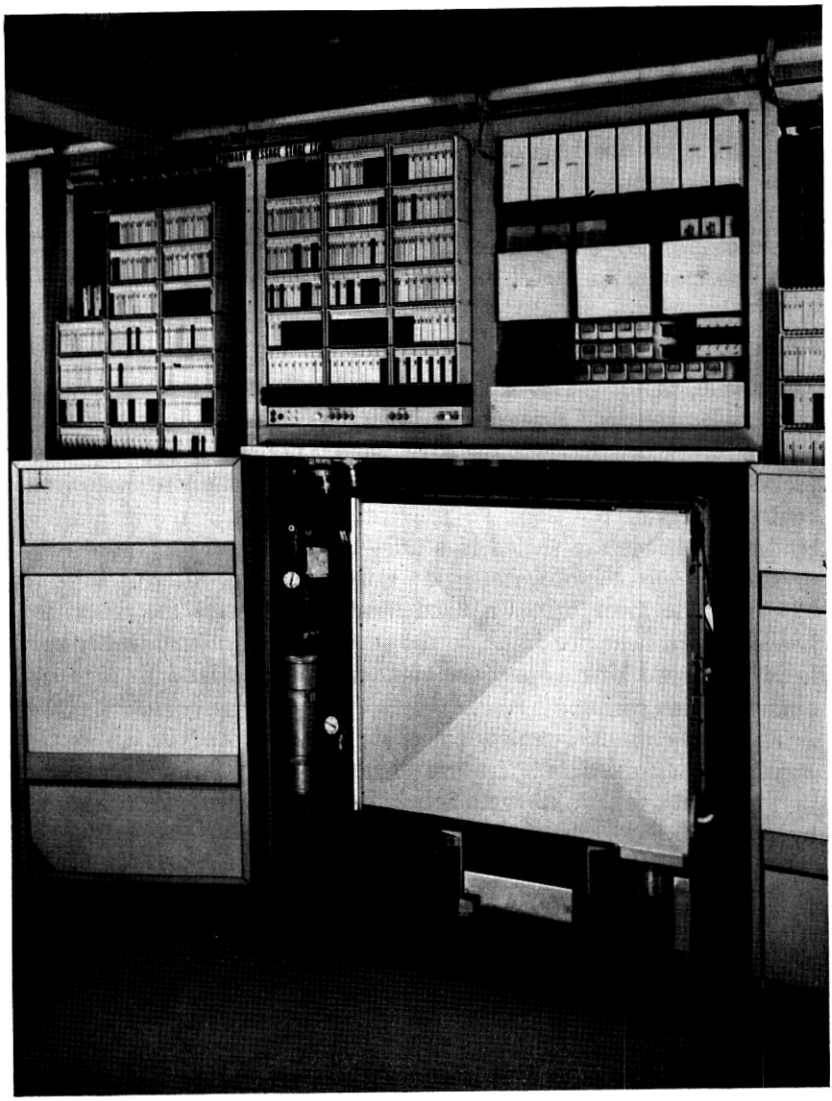

Fig. 2—Message store.

2.1 *Nitrogen System*

The disk heads are flown close to the memory surface by a closed pneumatic gas system. The gas is provided from a bottled nitrogen system mounted remote from the logic frame. The nitrogen system is duplicated with each simplex system, providing working pressure to a simplex half of a duplicated disk. Thus, a single failure in the nitrogen system will not cause a complete loss of disk facilities in either community. The pressure is reduced from the bottle pressure to the required 40 psi by a series of three regulators. The use of multiple regulators provides over-pressure protection to the file. Monitor alarms are provided to allow program detection of out-of-limits gas pressure.

III. SYSTEM INTERFACE

The central processor communicates with the message store via the high-speed buffer control call store memory and a wired logic buffer control (Fig. 3). This arrangement allows the processor to transfer

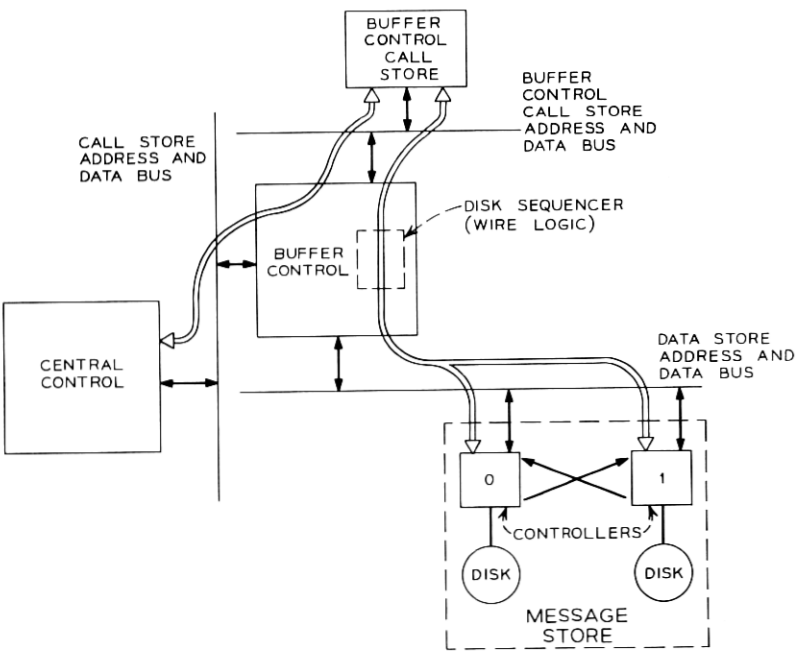

Fig. 3—No. 1 ESS ADF message store system interface.

disk data to the buffer control call store at a 5.5 microsecond/word rate. The wired logic sequencer, in the buffer control, then transfers the disk data from the buffer control call store to the primary message store at the 80 microsecond/word disk rate. Thus, the processor is released to continue its real-time work and not required to scan the message store for its asynchronous requests on a real-time basis.

The wired logic disk sequencer is one of the four peripheral sequencers, centrally located in the buffer control, which time-shares the common buffer control call store memory. A wired logic bus sequencer handles all requests for use of the buffer control call store bus by the common peripheral sequencers, and grants priority in a predetermined order.

The disk speed and phase are asynchronous to the fixed buffer control processing cycle and require the message store to issue commands when it is prepared to read or write the disk. The following requests are sent to the buffer control disk sequencer by private requests leads:

(*i*) *Instruction Request*—Message store ready to receive an instruction from buffer control for the current sector.

(*ii*) *Data Request*—Message store ready to receive a 24-bit data word (write mode) or requesting its 24-bit data register to be read (read mode) by buffer control.

(*iii*) *Status Request*—Message store requesting the buffer control to read its error status register to determine error history over the last block of 32 data words.

These requests are generated in the primary message store by decoding the *word* clock which is permanently magnetized on the disk surface. A *bit* clock provides a train of pulses to serially shift the contents of the 24-bit primary message store data register, one bit at a time, onto the disk.

A maximum of 12,288 bits of data (16 sectors/revolution by 32 words/sector by 24 bits/word) can be written or read during one revolution of the disk (40 milliseconds).

IV. DATA ORGANIZATION

To maintain a nearly uniform packing density of data on the disk surface, a disk is divided into three zones, as shown in Fig. 1. Each zone has its own clock-determined bit rate. Zone 1 has a bit frequency of 1.1 megabits/second, zone 2 has a bit frequency of 1.5 megabits/sec-

ond, and zone 3 has a bit frequency of 1.9 megabits/second. These three frequencies provide a nominal packing density of data on the disk surface of approximately 1,100 bits per inch. The disk surfaces are further segmented into 16 sectors. Each of the 16 sectors is composed of a *control* sector followed by a *data* sector.

The data portion of the sector is a block of 32 words written in an interlaced arrangement, as shown in Fig. 4. Zone 1 has three blocks interlaced, zone 2 has four blocks interlaced, and zone 3 has five blocks interlaced. Hence, the first, fourth, seventh, etc., words form a block of data in zone 1. Only one interlaced block may be written or read during a given sector and zone. This interlacing of data provides a sufficiently low data rate on the data store bus so that it may be shared simultaneously among many units. It also allows the duplicated and synchronized message store to read and write data at slightly different times and then, by individual time buffering, perform bus operations simultaneously.

4.1 *Clock and Timing Generation*

Each zone has a *bit* clock and a *word* clock used in reading and

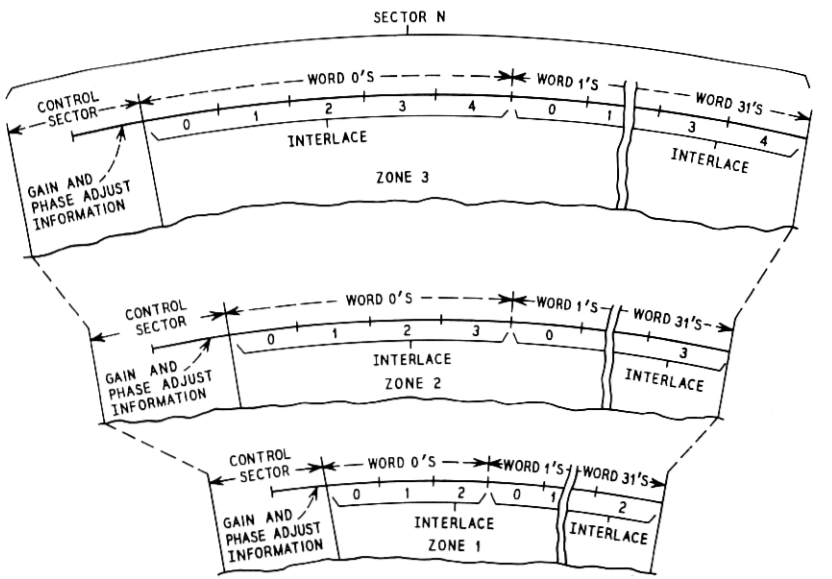

Fig. 4—Disk control and interlace arrangement as determined from word clock.

writing of data in that zone. The bit clock defines the bit positions within a word. The word clock provides word, interlace, and miscellaneous control information. The information in the word clock is coded in a binary, serial manner. This serial bit stream is shifted through a 9-bit shift register and periodically decoded with *decision time* pulses to obtain the disk logic timing (see Fig. 5). The decision time pulse is derived from a logical combination of the bit and the word clock. Normally, every pulse in the word clock has a corresponding pulse in the bit clock. Periodically, however, a pulse appears in the word clock with no corresponding pulse in the bit clock. This condition generates the decision time pulse, which causes decoding of the present contents of the shift register. The disk word clock generates ten timing marks in this manner.

The timing marks thus generated define interlace boundaries, start of sector, end of sector, and necessary timing pulses in the control sector for communications with the processor. In addition to generating the timing mark at a decision time, some timing marks cause gating

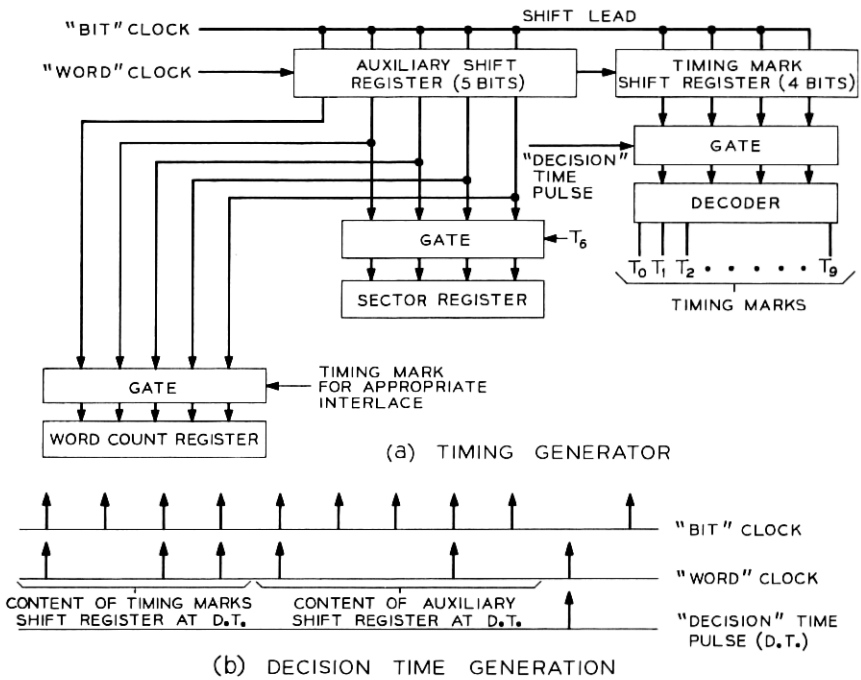

Fig. 5—Relationship of disk word and bit clocks.

of information in the auxiliary shift register to special purpose registers (see Fig. 5). In this manner positional information such as sector and word number is available without counting time marks. The clock information is written in the clock tracks by a special apparatus previous to the installation of the disk file into the message store. Once there, the clocks are write protected and can only be read by the disk logic circuits.

V. SERVO SYSTEM

The disk stores are operated in pairs, so if one should fail, the other will carry on. The two disks are servoed to permit matching of their outputs while they are being read and to permit simultaneous storage of messages.

Each disk reads its sector location from the *word* clock track. The 4-bit sector address is cross-coupled to its duplicated mate, and each disk compares its address with its mate's address. One will be slow, the other fast. The sign of the error controls the frequency of the servo oscillator, causing the disk to speed up or to slow down. The speed is controlled by a 360-Hz voltage-controlled oscillator. The oscillator is counted down to control the three-phase 60-Hz power for the motor.

All power for the message store is derived from the office battery. The disk motor operates from 208 volts, 3-phase, 60-Hz. A transistorized inverter, controlled by the counted down 360-Hz oscillator, inverts the battery voltage to the 208 volts required for the motor. The motor rotates nominally at 1725 rpm and is belt-coupled to the disk shaft. The disk speed is reduced to 1500 rpm by the belt pulley combination.

The following is a list of some of the pertinent design requirements which the *servo system* must possess:

(*i*) fully automatic,
(*ii*) no master-slave relationship,
(*iii*) all servo positional and frequency data derived from the normal clock signals, and
(*iv*) four-microsecond maximum phase error between servoed disks. (Forty milliseconds per revolution.)

Since a small phase error is required, a high phase gain is necessary. To make the system stable with high gains, rate feedback is required. The rate feedback could be derived by comparing the speed of two disks; instead, it is derived by comparing the disk clock to an internal

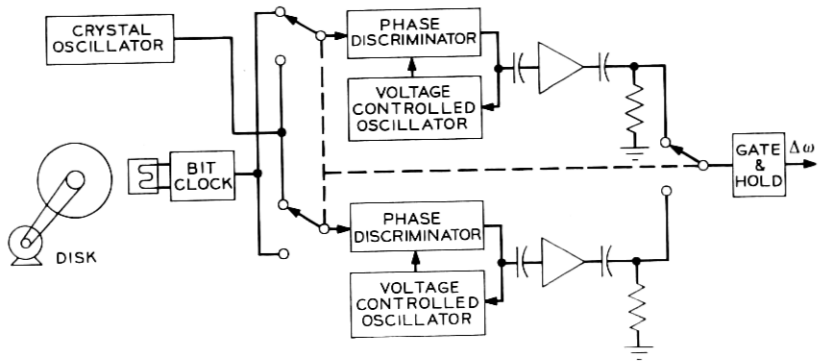

Fig. 6—Frequency comparator.

crystal oscillator. This eliminates cross-coupling frequency information and causes each disk to rotate with the same accuracy as its crystal. The frequency servo is shown in Fig. 6.

There are two paths through the servo circuit, phase and frequency. The phase error, leading or lagging, is derived from the comparison of the local disk address with the remote disk address. The result of this comparison is stored in a flip-flop which, if set, speeds up the oscillator, and if reset, slows down the oscillator. The oscillator controls the speed of the motor, hence, that of the disk. The position of the disk is repeatedly read and compared with the other disk, thus closing the loop (see Fig. 7).

In the second loop, the *bit* clock on the disk is compared with a reference crystal. The comparison results in an output which will speed up the oscillator if the bit clock is at a lower frequency than the crystal or will slow down the oscillator if the bit clock is at a higher frequency.

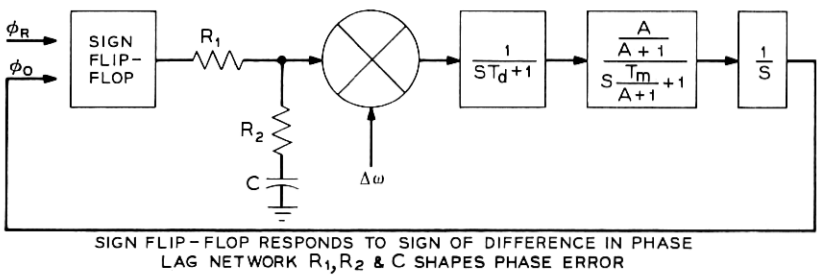

Fig. 7—Complete sign only servo.

Again, the oscillator output controls the speed of the motor, therefore, the frequency of the bit clock. The frequency loop always opposes the phase loop or any other force trying to change the speed of the disk.

Considering initial response, the oscillator, inverter, and disk can be represented by a simple time constant T_m.

$$T_m = \frac{I_m}{K_m} \quad \begin{matrix}\text{(total system inertia)} \\ \text{(slope torque-slip curve).}\end{matrix}$$

Figure 8 shows a servo structure without rate feedback. This simple servo would not meet the phase accuracy objectives.

To meet the phase accuracy objective, the phase loop gain must be high, which will cause the circuit to oscillate. Rate feedback is added to stabilize the circuit but must also meet severe requirements. Some of the problems are a result of the methods of powering the disks. Each phase of the power inverter produces a square wave output and is Y connected to produce the 3-phase power for the motor. Because the oscillator frequency is not changed until a transition occurs in the square wave output, a 0 to 2.77-ms delay is introduced into the loop. There are other small delays which are lumped into one time constant T_d.

Figure 9 shows rate or frequency feedback added to the servo loop illustrated in Fig. 8. For $A \gg 1$ and $T_m \gg T_d$ and $T_m/A \gg T_d$, an appropriate equivalent circuit can be derived. Starting with a 2-pole system, the closed loop response is:

$$(F_s)/(1 + F_s) = (A/[(ST_m + 1)(ST_d + 1)]) \div (1 + [A/(ST_m + 1)(ST_d + 1)])$$

$$(F_s)/(1 + F_s) \approx (1/\{[(ST_m)/A] + 1\}) \times (1/(ST_d + 1)).$$

Figure 10 shows the equivalent circuit as suggested by these equations.

Instead of proportional phase error, the actual servo uses sign only feedback, equivalent to infinite gain. For the slightest possible phase error, full output is obtained from the sign flip-flop. To reduce the frequency loop gain to a stable level, a lag network is added in the phase loop. The final servo design is shown in Fig. 11. Having infinite gain, the servo will continually oscillate with a predictable small phase error.

The frequency loop has two severe requirements: it must introduce

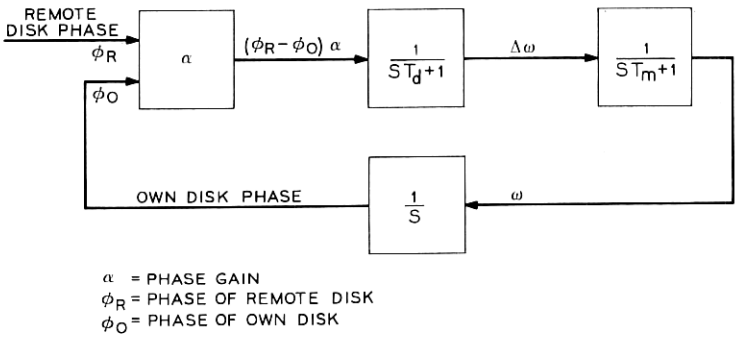

Fig. 8—Servo without rate feedback.

as small an open loop offset as possible, and it must derive the rate feedback with as small a delay as possible. For a gain A, of 100 and $T_d > 7.5$ ms, the frequency loop alone is unstable. T_d includes the 0 to 2.77-ms delay caused by the square wave drive to the motor. The requirements are an offset less than one part in 2×10^4, and effective delay of one millisecond or less, and a gain of 100.

The disk clock (1.1 MHz) is compared to a crystal oscillator. In one millisecond there are 1100 transitions, but the required accuracy is one part in 20,000. Thus, a counting type measurement of the relative frequency of the disk and clock is not practical. The comparison circuit must give high accuracy, low offset, and be immune to noise or variation in bit clock timing of single bits. The circuit developed uses a phase-locked oscillator as a frequency measuring device. An

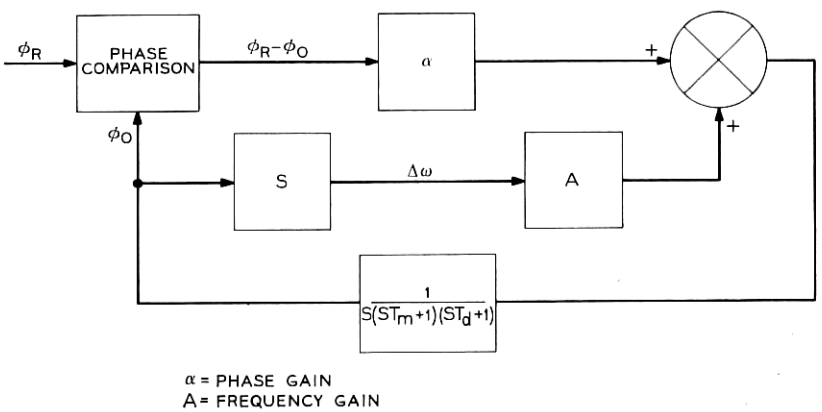

Fig. 9—Servo with rate feedback.

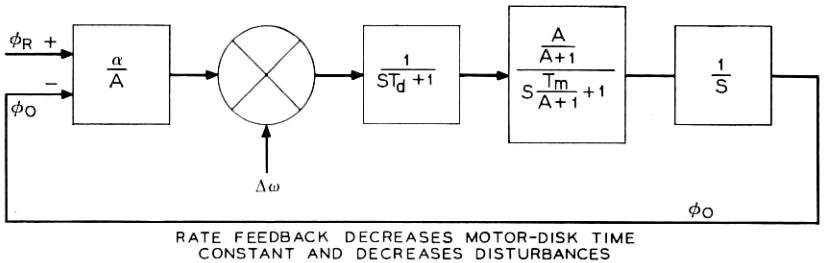

Fig. 10—Effect of rate feedback on servo.

RC oscillator is voltage controlled by the output of a phase discriminator.

This voltage-controlled oscillator is phase locked first to the local crystal, then to the disk clock. The change in the oscillator control voltage is proportional to the difference in frequency of the disk clock and the crystal. Since only the change in voltage is used, there is no offset error caused by changes in the voltage-controlled oscillator components or supply voltages. The change in voltage is amplified and rectified by a phase sensitive circuit. To provide a continuous output proportional to disk speed, two circuits are used. One circuit is locked to the crystal oscillator while the other is locked to the disk clock, thus producing the desired output.

Figure 6 is a simplified schematic of the frequency compensation circuit. The voltage-controlled oscillator can have any phase relationship when switched from crystal to disk or vice-versa. This large phase error appears as a pulse on the control voltage lead. This pulse is gated from a holding circuit so it cannot upset the actual error signal derived from the control voltage.

Figure 11 shows a two-disk system and the servo circuit. When first turned on, the two disks may have any phase relationship and differing speeds. For large phase errors, the error voltage is increased to decrease the time necessary to approach the same phase. Large phase error signals are derived from comparison of the more significant bits of the address. When the phase error is small, the comparison is switched to least significant bits, and the error voltage is decreased. All sequencing is automatic.

When the disks are servoed, the final speed is proportional to the average of the two crystal frequencies if there were no other offsets. Since there are other offsets, the final frequency is near the crystal

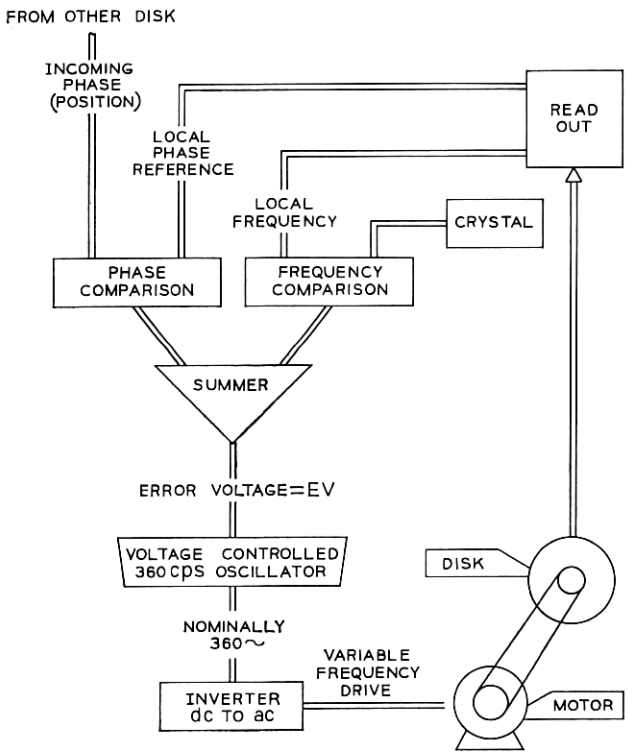

Fig. 11—Basic servo system.

frequencies but is not necessarily the average. If one servo stops working, and the disk continues to rotate, the other disk will stay phase locked to the unservoed disk. By deriving the phase error from the sector address and the rate feedback from the bit clock, the disk is affected with the same clock requirements as it would be without the servo.

The disks have been servoed for periods of months with continuous recording of phase errors with no error in excess of four microseconds.

VI. READ/WRITE CIRCUITS

Accessing a track to read or write is accomplished by selecting a head in the 96 by 12 head matrix (Fig. 12). Figure 13 shows a typical head group and its associated selection circuitry. In the writing mode, the emitter of Q2 is raised +12 volts to select the desired center tap.

A current regulator, Q5, is switched on, causing 130 milliamperes to flow in the ground clamp, Q6. Then the write switches for the select head group are enabled to write the data pattern.

The recording method used on the disk file is nonreturn-to-zero, where the disk surface is magnetically saturated in one direction for a "1" and the other direction for a "0". This recording method is schematically shown in Fig. 14. If switches Q3 and Q4 of Fig. 13 are alternately switched on and off, the 130 milliamperes will flow in one side of the head or the other. The fringing field in the gap region of the head (Fig. 14) will magnetize the disk surface in relation to the direction of current flow in the head windings.

While signals are being read from the disk surface, the write switches and current regulator are turned off. One of the 96 center taps is returned to +12 volts through transistor Q2, and read switch, and transistor Q7 is grounded, to access the desired head group. This causes a small current to flow from the center tap of a select head through the 1K ohms read-select resistors to ground. This forward biases all diodes between the desired head and the read amplifier, providing the low impedance path for the millivolt head signals to reach the amplifier. Because of the heads' poor low-frequency response, the read signals are differentiated pulses whose peaks correspond to the magnetic transitions on the disk surface. These signals are amplified and detected to obtain the nonreturn-to-zero representation.

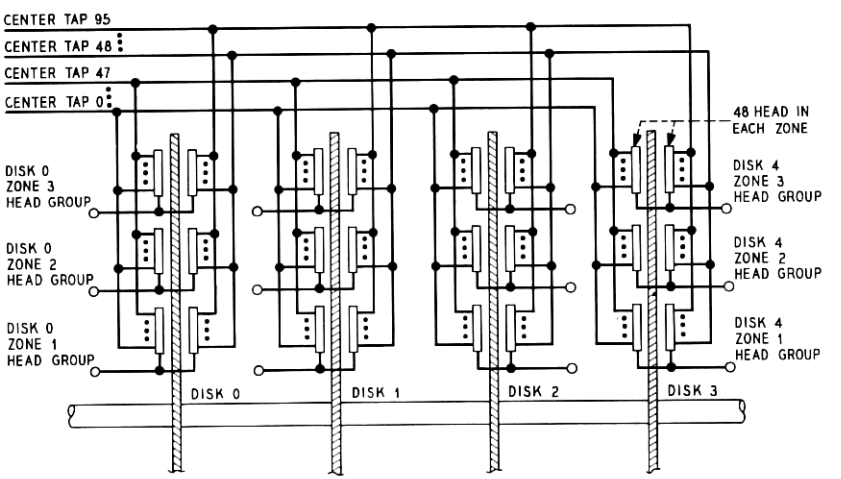

Fig. 12—Data track selection in a fixed head/track disk store.

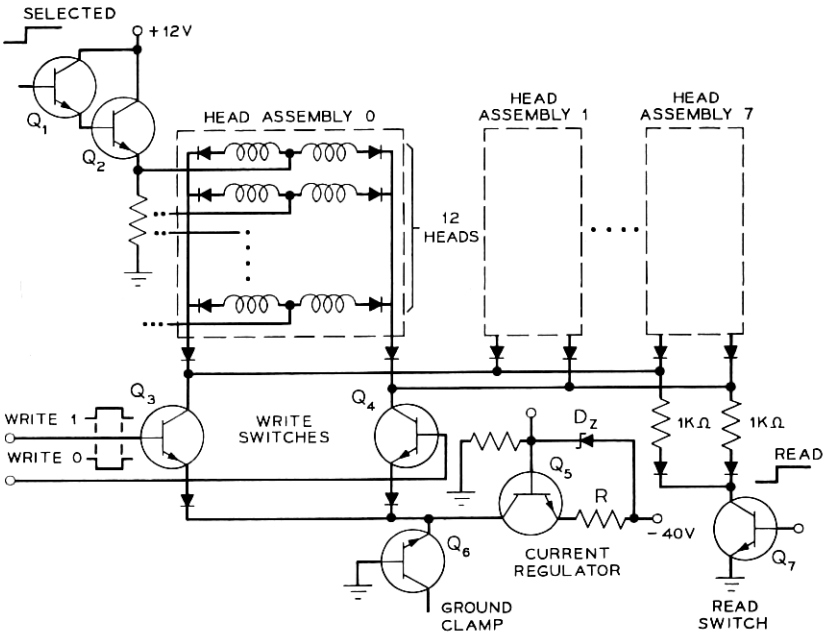

Fig. 13—Head selection read/write.

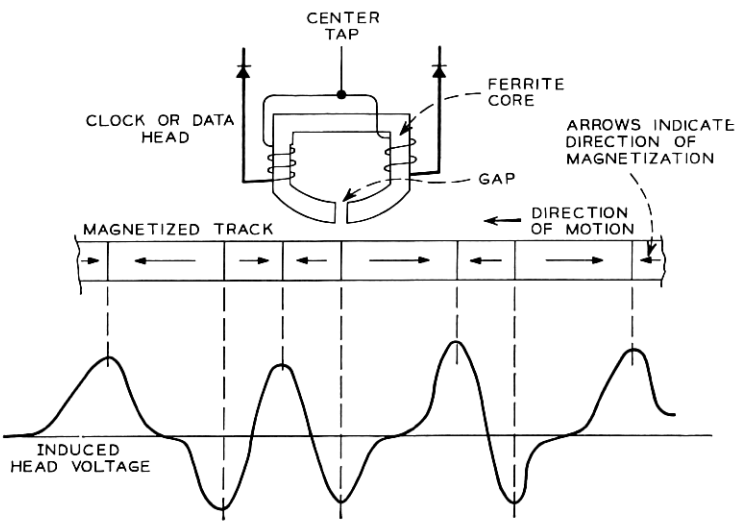

Fig. 14—NRZ recording.

6.1 *Disk Readout Detector*

In the readout detector of Fig. 15a, the analog signals are converted to a digital pattern similar to the nonreturn-to-zero writing current. Because peaks of the analog signal most accurately define the flux transitions of the disk surface, the peak detector is an important part of the detector. When the voltage level of the signal exceeds a fixed threshold, the threshold detector's output goes to a "1" state. When the peak detector senses the peak, its output becomes "1" and provides the necessary conditions to activate the following NAND gate. The output of the NAND gate then sets or resets the flip-flop depending upon which output is being considered. Hence, the data

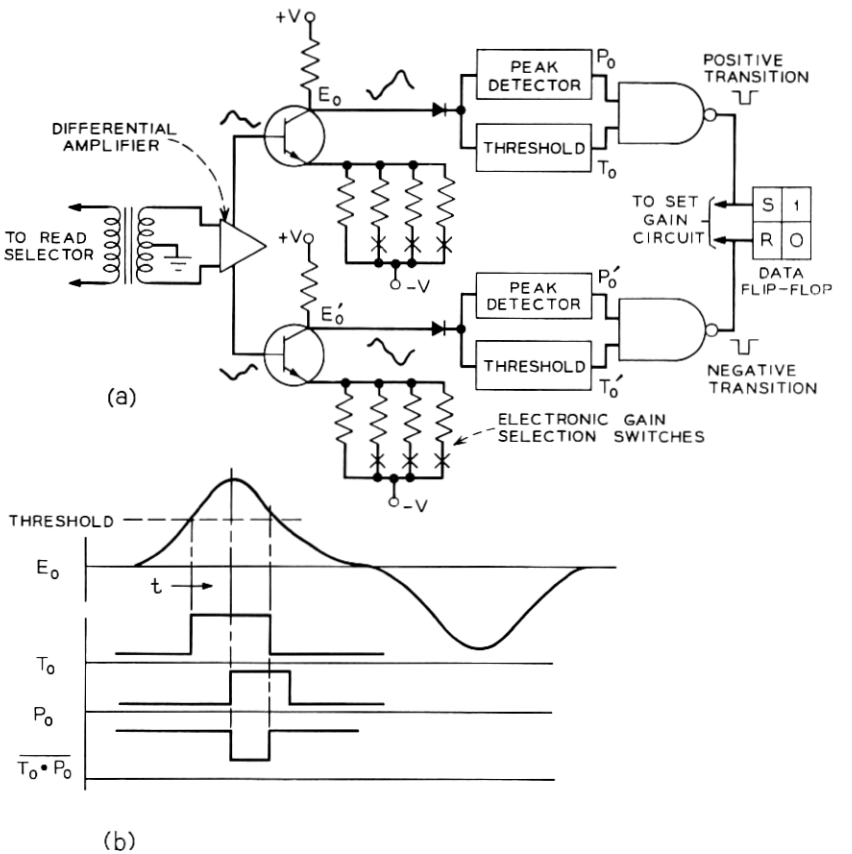

Fig. 15—Readout detection scheme: (a) detector circuits, and (b) detector circuit wave forms.

flip-flop is set for a positive peak and reset for a negative peak (see Fig. 15b for timing). The resultant output of the data cell is a non-return-to-zero representation of the data on the disk surface. This output must then be shifted, bit by bit, into the data shift register with the *bit* clock until an entire serial word is assembled and is available in parallel form.

6.2 *Adaptive Gain Circuit*

To accommodate the 5-to-1 signal amplitude variations from the disk heads, an adaptive gain amplifier is used in the readout circuit. The final stage of the readout amplifier incorporates a fixed collector resistance with a switchable emitter resistance which can vary the final stage gain from 2.5 to 10.7.

In the *control* portion of each sector on every track, there are 21 permanent check bits (see Fig. 16). These bits characterize the signal in the following data portion of that sector. They provide the phase and gain information necessary to read that sector track address accurately. The first two transitions of the 21 check bits are used in setting the gain of the amplifier. Approximately 60 microseconds preceding the gain check bits, the head access circuitry selects the desired head for the sector to be read. The 60-microsecond period is required for the amplifier to recover from the large head switching transient. The initial amplifier gain setting is such that the amplifier

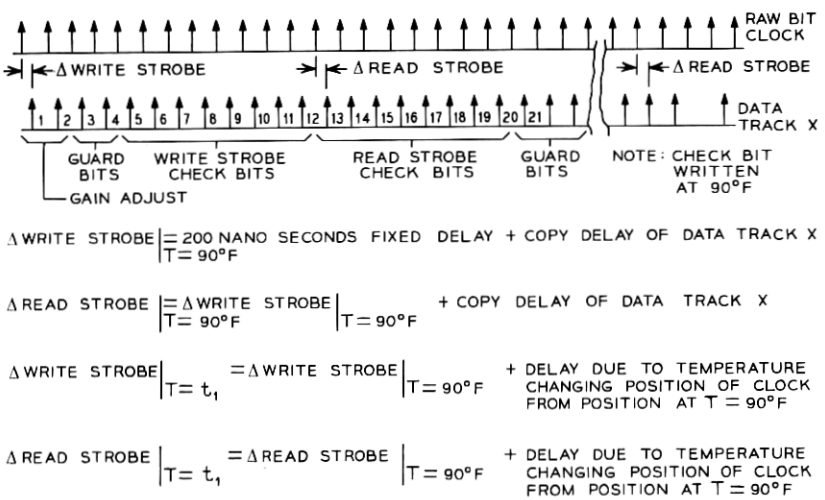

Fig. 16—Adaptive write/read strobe timing relationship.

side used in detecting negative peaks has a final stage gain of 2.5, and the other side has a gain of 3.7. While the two check bits are passing through the detector, both the positive transition and the negative transition outputs are monitored (see Fig. 15a). If the first check bit causes the negative transition output to pulse, the gain of both sides of the amplifier will be set to a low gain of 4.8. If only the positive transition output pulses as a result of the second check bit, an intermediate gain of 7 will be set on both sides of the amplifier. If neither output pulses, a maximum gain of 10.7 is selected. Hence, the entire adaptive gain selection process is performed with only one gain change of the amplifier.

6.3 *Adaptive Data Strobe Circuit*

The remaining 19 check bits in the control sector characterize the unique write-to-read delays of each head-track combination. They further compensate for the changes in the clock and data timing relationships with temperature excursions. The clock heads and data heads are mounted on a large, aluminum casting. Because the two types of heads are located some distance apart, the temperature distortions in the casting move the data with respect to the reference bit clock. If the data is written with respect to the bit clock at the lowest ambient temperature and read back at this temperature, only a fixed delay need be added to the clock to obtain strobing pulses to read out the data. However, if the temperature is raised, distortion of the casting may move the data with respect to the clock; this total movement can exceed a data bit time. Hence, the data written at one temperature extreme cannot be retrieved at the other temperature extreme.

To accommodate the temperature range required in the No. 1 ESS ADF, two sets of check bits are written in each control sector before a disk is put into service. The first set of check bits (eight bits) is written with a fixed delay of 200 nanoseconds at a nominal temperature of 90°F. Due to circuit delays, this fixed 200- nanosecond delay places the first set of check bits roughly 400 nanoseconds away from the clock, as seen at the readout (see Fig. 16). However, this is true only at the check bit writing temperature. After the first check bits are written, they are read and variable delay is added to the clock in an amount necessary to equal the write-to-read delay plus the fixed 200-nanosecond writing delay. This places the delayed clock in alignment with the first set of check bits. With this delay in the clock, the

second set of check bits is written directly behind the first. These delay times are described in Fig. 16. After check bits have been written on all tracks, the disk is ready to be placed into service. Once written, the two sets of check bits are permanent and are rewritten only if destroyed or if the replacement of a circuit pack in the read/write path affects the write-to-read delay.

A method similar to the writing of the second set of check bits is used to write into a data sector. Upon receiving the write command, the strobe circuit reads the first set of check bits. It places a delay in the clock equal to the time difference between the first set of check bits and the reference bit clock. This delayed clock, called the write strobe, is then used to write the data in an interlace specified in the address register. Hence, the effect of adjusting the write strobe clock is to place the data in the same position relative to the first set of check bits, independent of the phase shifts in the reference bit clock with respect to the first set of check bits. Reading previously written data, the strobe circuit reads the second set of check bits, which are in the same relative position to the clock as the data, and chooses an optimum read strobe clock for reading the data. Therefore, the timing used in writing and reading of the data on a particular track depends only on phase information contained in the same track as the data and not on a phase varying clock located on a remote track. The range of the strobe circuit (± 450 nanoseconds) accommodates a maximum temperature shift of $\pm 30°F$ within the disk file enclosure.

Another advantage of the write strobe technique is that one set of check bits suffices for all interlaces in a track. Also, this technique requires only a minimum guard gap between each interlace because, independent of temperature, the data in an interlace is always written in the same position on the disk surface.

VII. DATA TRANSFER

A message store administration program, resident in the central processor, loads a disk instruction and a data block address into dedicated 2-word task hoppers in the buffer control call store (Fig. 17a, 17b). Data being transferred to disk is stored in a 32-word data block (Fig. 17c) beginning at the address specified by the contents of the buffer control call store data block address. No further action by the processor is required to carry out the disk data transfer.

The wired logic buffer control disk sequencer bids for and is granted use of the buffer control call store facility for one cycle. The sequencer

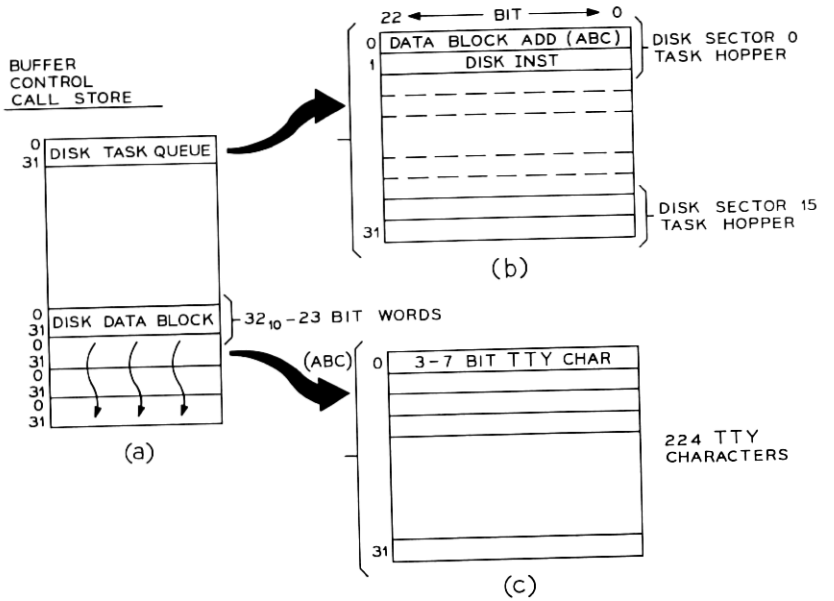

Fig. 17—Message store task queue and data block layout.

reads the 2-word task hopper (address specified by a wired machine constant plus the 4-bit queue counter) and transfers its contents to the disk sequencer address and data registers (Fig. 18). The first hopper word (buffer control call store data block address) is gated to the sequencer address register, the second (disk instruction) is gated to the sequencer data register. Once this preloading is completed, the disk sequencer waits for an instruction request from the message store. After receiving the instruction request and being granted permission to use the data store bus, the sequencer gates the contents of its data register (disk instruction) onto the data store write bus and into the message store address register. This instruction is decoded by the message store to select a unique data head and clock for transferring data.

The disk sequencer then preloads the first data word into its data register (disk write) and waits for the asynchronous data request from the message store. The address of the data stored in the buffer control call store is determined from the disk sequencer address register. When the data request is received by the disk sequencer, the contents of the sequencer data register (disk data) are gated onto the

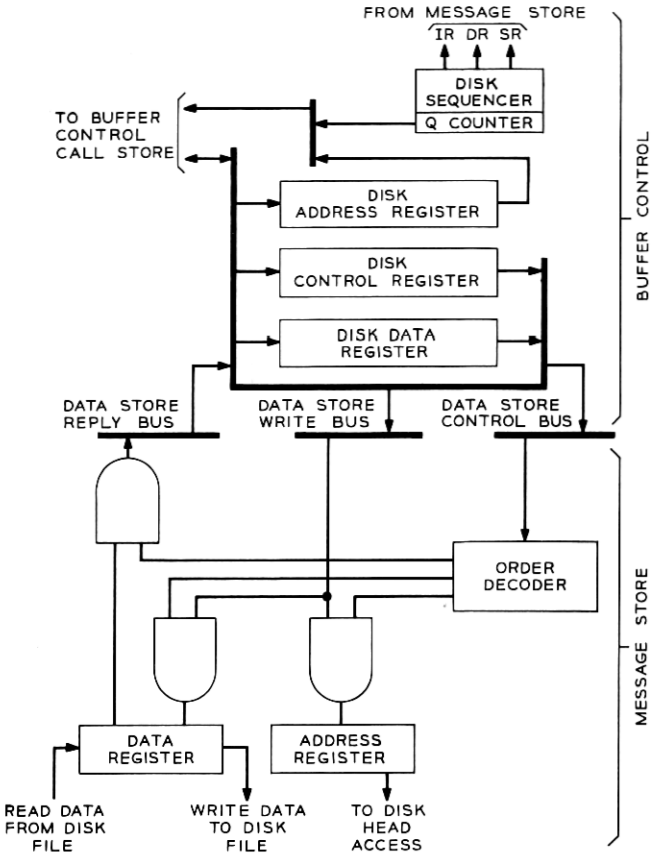

Fig. 18—Wired logic sequencer register access.

bus and into the message store data register. From there, the 24-bit word is shifted, one bit at a time, onto the disk memory surface. This sequence is repeated 32 times to complete the transfer of the 32-word buffer control call store data block into a disk sector. Reading disk data is accomplished in a similar fashion.

Upon completion of the block transfer, the disk sequencer waits for the asynchronous status request from the message store. Once the status request is detected, the sequencer reads the contents of the message store status register (Fig. 19) into its data register. If the status report indicates success, a no-op (all ones code) is written in the buffer control call store task hopper to indicate successful comple-

tion of the transfer. If a failure is detected by the message store or disk sequencer, the entire operation is left in the hopper for one additional retry during the next disk revolution. Failure on the retry causes the buffer control call store task queue instruction to be overwritten with a failure status report. The failure status report serves a dual function. First, it alerts the client program requesting the disk data transfer that the transfer was unsuccessful. It is also used as an input to a maintenance fault recognition program which must isolate the faulty unit and prepare it for programmed diagnostic testing.

7.1 *Synchronization*

Sector synchronization must be maintained between the message store and the wired logic disk sequencer. Thus, when the disk is prepared to process sector n (determined by the disk position), the sequencer is synchronized to read the current instruction for sector n from the buffer control call store task queue. The sequencer contains a 4-bit binary queue counter (0 through 15 disk sectors) synchronized to the disk sector. Each status report read from the message store specifies the 4-bit sector address of the disk sector just transferred. To synchronize the queue counter to the current message store sector, the sequencer is programmed into a startup mode. When the first status report is read by the sequencer, the 4-bit message store sector address is extracted and written into the queue counter. The queue counter is then used to read the proper task hoppers associated with the nth sector. The counter is incremented by the status request from the message store from 0 to 15 and recycles to 0 at the start of each new job cycle. Once synchronization is achieved, no further initialization is required.

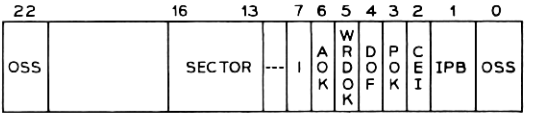

OSS = OVERALL STATUS SUMMARY
IPB = I AM PROBABLY BAD
CEI = COMMUNICATION ERROR INTERNAL
POK = DATA PARITY FAILURE
DOF = DATA OVERFLOW
WRDOK = WORD OK
AOK = ACCESS OK

Fig. 19—Status report exchanged between message store and wired sequencer.

VIII. READ/WRITE FAULT DETECTION

A large portion of the message store circuitry detects faults in the active read/write circuits. In the read mode, with normal system configuration, matching data with the duplicate disk controller provides a powerful fault detection tool. In addition, every 24-bit word has an odd parity bit which is checked as the data is read serially from the disk. When an entire word plus parity is in the 24-bit data shift register, the register is tested for at least a one. This check, plus parity, will detect most faults in the readout if the disk communities are not running in a duplicate matching mode. It also will help identify the faulty unit when a mismatch occurs in the match mode.

The write mode is more thoroughly monitored because, once the data leaves the wired logic sequencer in the buffer control, it is no longer matched. Each of the message store frames writes its data independently. The buffer control sends odd parity with each word to be written on the disk; as data is shifted onto the disk, parity is checked. After the 24-bit word is shifted out of the register, it is tested for all zeros. In addition, for the 2-bit times preceding each word, the write current is turned on and tested for 130 milliamperes ± 10 percent. Also, the head selection matrix is tested for exactly one center tap selection and one group selection. Having passed all the above tests, the probability of writing bad data on the disk is minimized.

IX. MAINTENANCE

The message store includes maintenance hardware to check for the proper responses to read/write commands, and to verify proper data handling. This hardware, along with maintenance programs located within the central processor, alert the system to all irregularities encountered during disk data storage or retrieval. The maintenance functions are as follows:

(*i*) Hardware checks are performed by the message store, and failures cause system interrupts (immediate action).

(*ii*) Hardware checks are performed by the message store, and failures cause bad status reports to be written in the task queue (deferred action).

(*iii*) Routine programmed diagnostic tests are run to exercise all maintenance associated hardware.

All data transfers between the buffer control call store and the mes-

sage store are protected by an odd parity bit and an acknowledgment (all-seems-well) signal. Bad parity or the lack of an all-seems-well response will cause an immediate maintenance interrupt in the central processor (Fig. 20). Maintenance interrupts also occur when data read from disk by the duplicated sequencer fails to match.

A maintenance interrupt detected in the central processor stops normal program flow and transfers to an interrupt filter program. Before proceeding to isolate the faulty unit, all processor registers are saved so that the normal program flow can be reentered once the faulty unit is removed. The filter program interrogates unique hardware error indicators to determine the failing subsystem. Control is then transferred to the fault recognition test program for the failing subsystem. The failing unit is removed from service and its simplex mate configured to respond for both itself and the faulty unit. Once the faulty unit has been removed, and an automatic diagnostic test requested, normal processing is resumed at the interrupted program. Diagnostic requests are honored routinely and are run in a segmented mode with normal processing. Diagnostic test programs resolve the fault to within a few replaceable plug-in circuit boards and print out a trouble number associated with this location. After the faulty circuit pack is replaced, a diagnostic test is run to verify the subsystem. The message store is then updated from the active unit and returned to normal service.

The second maintenance function causes no immediate interrupts. A status report is read by the wired logic sequencer at the end of each disk sector. The message store replies with a data word indicating the results of hardware checks performed during the previous block transfer. If the sequencer detects a status report failure, and the data transfer is a second attempt, it overwrites the corresponding instruction word in the buffer control call store task hopper with the failure report. Thus, a failing data transfer is automatically abandoned by the hardware, and notification given to the client administration program by the failure status report in the task queue. The presence of a failure status report in the task queue causes the administration program to enter the fault recognition program at the block repeat entry (see Fig. 20). Tests performed at the fault recognition level are sufficient to isolate and remove the faulty message store half and request diagnostic testing.

Maintenance circuitry is tested by routine daily diagnostics performed on the message store. The diagnostic tests exercise all message store hardware and read/write circuits.

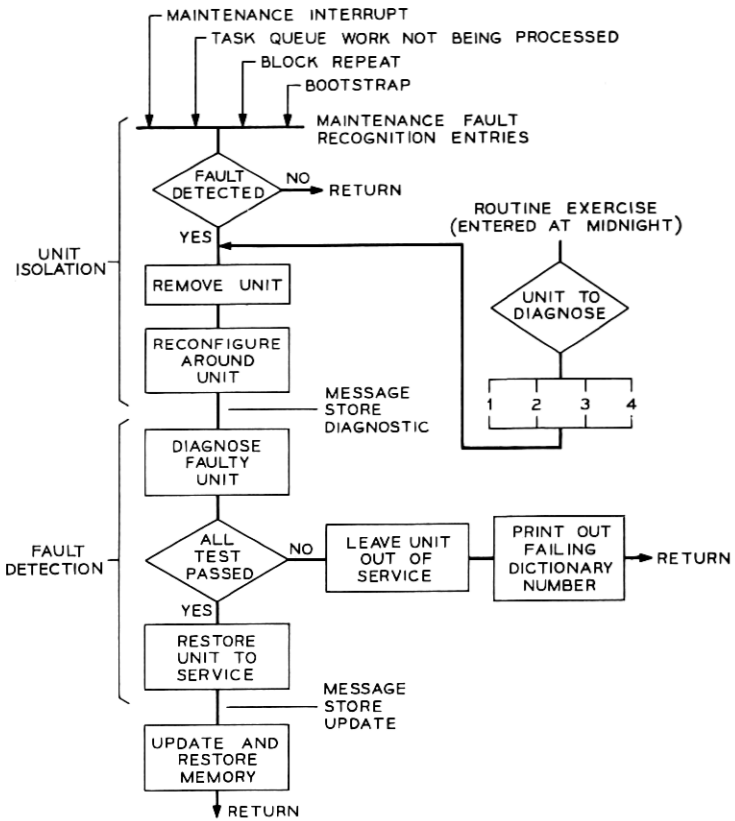

Fig. 20—Maintenance software flow diagram.

X. CONCLUSION

The message store is a 60-megabit disk file storage system. The cost per bit of this storage is two orders of magnitude less than the cost for the ferrite sheet call store employed in the No. 1 ESS. The disk memory system is run in full duplex to enable matching of all data by a duplicated wired logic sequencer. The integrity of the data is insured by internal hardware checks, all-seems-well, data parity, and cross matching of all bus transmissions.

Disk operation is asynchronous with the fixed 5.5-microsecond processing cycle. The disk cycle is determined by a set of programmed clocks permanently recorded on the surface. Two pairs of disks are held in frequency and phase synchronism by a closed loop servo.

The hardware is backed up by maintenance software programs to detect, isolate, and diagnose failing circuits. Operating experience, since system cutover on February 3, 1969, indicates that the primary message store meets expected reliability requirements. Faulty circuits have been detected and isolated by the field craftsman using a fault dictionary. Faulty data heads within the disk file have also been isolated by a special write/read failure matrix test provided as part of the diagnostic program.

XI. ACKNOWLEDGMENTS

The No. 1 ESS ADF disk memory was developed by many engineers, all of whom contributed significantly to the project. Mr. T. S. Greenwood led the disk development, and Mr. L. E. Gallaher supervised the disk liaison, read/write circuitry, and system design. Mr. P. K. Giloth headed the department which designed the operational administration software. The maintenance system was designed by Mr. R. W. Downing's Maintenance Department.

REFERENCES

1. Genke, R. M., Harding, P. A., and Staehler, R. E., "No. 1 ESS Call Store," B.S.T.J., *43*, No. 5 (September 1964), pp. 2147–2191.
2. Cagle, W. B., Menne, R. S., Skinner, R. S., Staehler, R. E., and Underwood, M. D., "No. 1 ESS Logic Circuits and Their Application to Design of the Central Control," B.S.T.J., *43*, No. 5 (September 1964), pp. 2055–2095.
3. Aitcheson, E. J., and Cook, R. F., "No. 1 ESS ADF: Maintenance Plan," B.S.T.J., this issue, pp. 2831–2856.