J. W. Irwin
J. V. Cassie
H. C. Oppeboen

The IBM 3803/3420 Magnetic Tape Subsystem

Abstract: The design innovations in the IBM 3803 Tape Control Unit and the IBM 3420 Magnetic Tape Unit are discussed. The new design concepts include a full readback parity check, a system of microprogram control that obtains high operating rates through overlap of channel control and motion control functions, a new phase-error detection system, a modified dual-density recording feature and a new interface between the control unit and tape drive. The combination of novel error detection methods, control features and programming support (including an expanded scheme of on-line diagnostics) contributes to greatly improved total system performance over previous models of tape subsystems.

Introduction

Although half-inch magnetic tape stores a great deal of information per unit volume at a relatively low cost, it does present certain limitations, the greatest being that of the time interval required for data access. To overcome this limitation, it would be necessary either to increase tape speed, increase density, decrease head-gap/data distance, increase the number of parallel tracks, or, for certain applications, raise the rewind speed. To maintain compatibility with earlier IBM tape subsystems, the development of the IBM 3803/3420 magnetic tape subsystem was concentrated around improved access time and higher rewind speed. The IBM 3420 is the new tape drive and the IBM 3803 its control unit.

Improved access time was achieved by positioning the readhead gap closer to the data, thereby reducing the access time interval in subsequent reads. Improved rewind time was achieved by more positive control over the tape as it enters the vacuum columns, and the control was obtained with a new configuration of tachometers for high-resolution tape speed information.

The IBM 3803/3420 is a total replacement for any existing IBM 2803/2400 tape subsystem. The former now includes features called a *two-channel switch*, which allows connection to two separate channels, and *device switching*, which allows up to four control units to communicate with up to 16 tape drives. These two features, available in all combinations, have existed only on certain models in the past, and the device switching had required a separate unit. The general configuration is shown in Fig. 1.

This paper discusses the following new design features of the IBM 3803/3420 tape subsystem:

Read-only storage control through microprogramming
Phase-error detection and correction
Full readback parity check
Radial interface for increased function and serviceability
Monolithic technology in the READ/WRITE channel
New approach to reel control
Greatly reduced dimensions of the control unit.

General description

A magnetic tape subsystem consists of an electronic tape control unit (designated in this paper by TCU), which connects one of several magnetic tape units (TU) to the computer's central processing unit (CPU). The TCU, primarily an electronic unit, controls the operation of the subsystem. Channel commands are received by the TCU and decoded so that the required command can be sent to the proper tape unit. When an operation is completed, the TCU responds to the channel with the appropriate indication of status.

When the processor is to perform a magnetic tape function, e.g., READ, WRITE or CONTROL, the command is sent to the TCU. If the operation is to be a WRITE, the TCU starts the tape motion and requests the first information from the processor. After all the information has been received, the TCU stops the tape motion and signals the processor that the operation is completed.

The performance goals of the IBM 3803 Control Unit introduced some unique design problems. Early in the

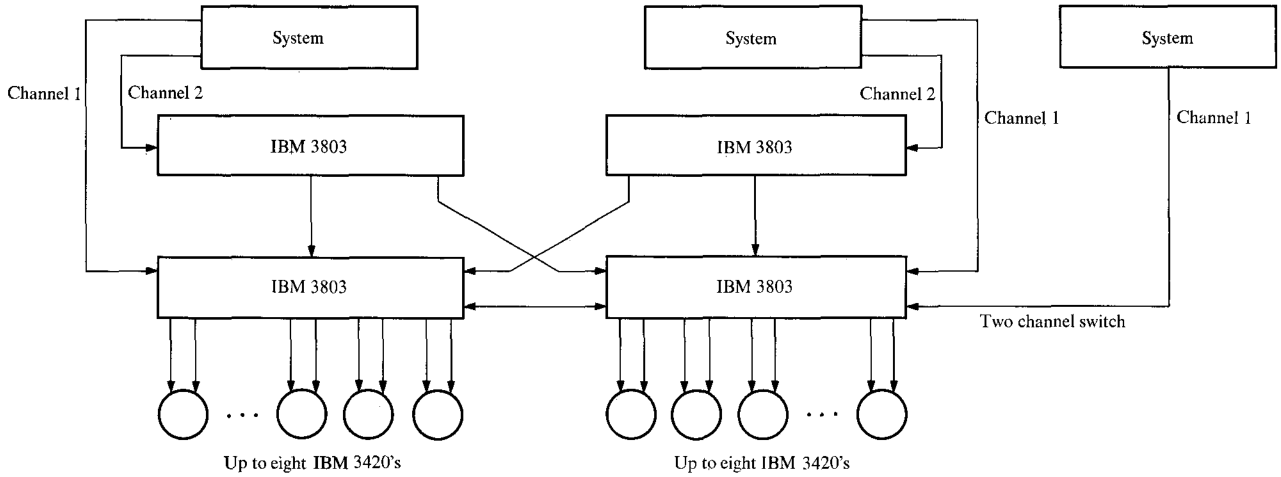

Figure 1 Configuration of the IBM 3803/3420 Magnetic Tape Subsystem showing how four control units communicate with up to 16 tape drives. This example of a typical configuration shows a pool of 3420 drives which can be accessed by three independent systems.

program it was decided to reduce the physical space required for the new subsystem to give the user the option of placing two IBM 3803's in the space previously occupied by one IBM 2803. The goal was to package the new TCU in a frame the same size as that of the IBM 3420 Magnetic Tape Unit. Another performance goal was to improve control unit design by the use of two separate microprogram units, which will be discussed in the following section.

• *Microprogram functions*

The 3803 design incorporates a microprogram which uses a read-only storage (ROS), a monolithic memory containing 1024 bits on a 16-pin module measuring 1/2-in. square. Monolithic memory chips are "personalized," i.e., chips are adapted for the specific purpose. In order to determine personalization patterns economically it was necessary to develop a READ/WRITE control store for early prototype machines. A monolithic READ/WRITE module was chosen which matched the speeds of the ROS used (50-nsec access and 100-nsec cycle time). The READ/WRITE memory was packaged in a separate instrument frame with its own power supply and a separate card reader for inputting the changing bit patterns used to control the 3803 TCU.

The ROS in the 3803 is mounted on logic cards plugged directly into the logic boards of the machine. The READ/WRITE control memory plugs into the same card sockets by means of signal cable connected to unpopulated ROS cards. This flexibility allows complete freedom to run the machine on ROS or to replace the ROS with the READ/WRITE memory if some microprogram changes must be verified prior to chip personalization. The present test results have shown this to be completely successful, and the machine operation is identical using either memory.

The 3803 has two identical microprogram units that operate separately. This innovation in control unit design was done for three reasons.

First, it allowed the control unit design to more closely parallel the traditional hard logic control unit design, where the control unit was divided into three basic sections of operation: the interface control section, the motion section and the read/write section. In the 3803 design the read/write section is retained in hard logic (although some control functions are contributed by the microprogrammed sections), and the interface control section and motion section are replaced by independent microprogrammed units. By maintaining the independence of these two machine sections, the control program of each section is considerably simplified. A single control program would have to perform the functions of both sections simultaneously. General functions are shown in Fig. 2.

Second, although it was known that the highest efficiency in the use of read-only store occurred with a relatively short microprogrammed control word, there were not sufficient decode functions in such a short microprogrammed word to perform all the functions required by a machine of this size. For instance, only a relatively small number of condition branches can be decoded from a 16-bit microprogram word. With two microprogram units the number of branch conditions is doubled.

The third reason for using two microprogram units is the ability to obtain high operating rates through overlap. A good example of such overlap is the initial selection sequence, where one microprogram unit decodes the command and responds to the channel while the second microprogram unit requests information on the tape drive status and assembles a composite picture of tape drive status and internal status remaining from previous opera-

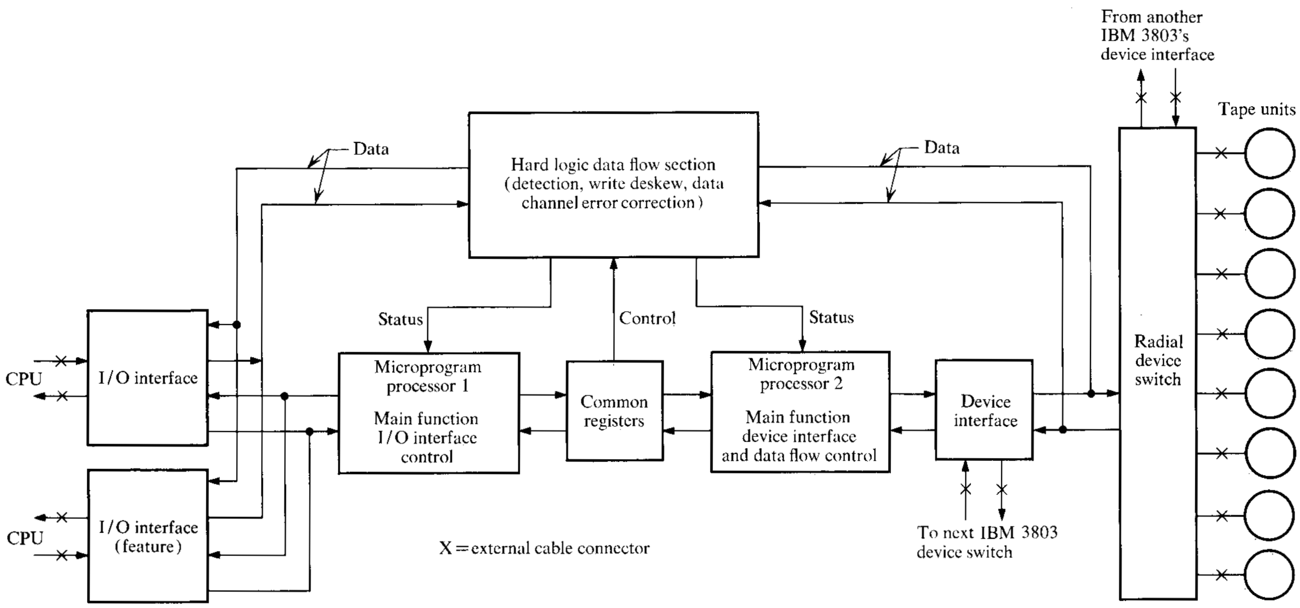

Figure 2 Block diagram of microprocessing.

tions. This information is then used to determine the proper response to the channel. Another important use of overlap is during READ and WRITE operations, where the first microprogram control unit monitors the channel activity while the second provides monitoring functions and control functions for the read/write section.

There is a master/slave relationship between the two microprogram units, in which the microprogram unit providing channel control functions is the master and the microprogram unit providing motion control functions is the slave. The master/slave relationship is enforced by a TRAP function. During TRAP, the operation of the microprogram unit is temporarily suspended and the instruction counters are reset to address zero. Proceeding from address zero a series of instructions is performed to branch the microprogram unit into the proper routine for the instruction in process.

During periods when the machine is not actively performing a command an important series of instructions is performed. First, the two microprogrammable units perform a closely interlocked scan sequence in which all attached tape units and all internal registers are scanned for error conditions and outstanding interrupts that must be presented to the channel. The synchronization between the two microprograms is provided by crossover registers and status registers that allow each microprogram unit to monitor the progress of the other. If any interrupt is found the microprogram units go into a WAIT loop until serviced by the channel. If there is no outstanding status each microprogrammable unit checks its own operation and then checks the operation of the other. The checkout routines test the functions of the arithmetic unit and all possible branch conditions and registers. If the checkout routines discover a malfunction, the control unit attempts to inform the channel of this fact (providing that sufficient operational capability still exists). The microprogram units continue to perform the interrupt scan and checkout routines until an interrupt or an error occurs, or until a subsequent operation is initiated. During the "power-on" sequence a considerably more extensive checkout routine is performed that includes test of all possible interface functions and interface branch conditions.

- *Master clock*

In order to perform all functions required of a microprogram cycle within 150 nsec it was necessary to construct a master clock capable of dividing that period into much smaller increments. In order to clock reliably at this rate it was necessary to construct the master clock from a circuit family that was somewhat faster than that used for the remainder of the machine. The master clock in the 3803 is operated by an oscillator at a rate in excess of 20 MHz and produces time divisions of <25 nsec.

The first 3803 engineering models used master clocks composed of delay lines, which were driven and constantly corrected by oscillators to assure the accuracy of microprogram timing loops. However, it was found impossible to control the overlap between successive timing pulses that lead to many difficulties within the microprogramming unit logic. In addition, although the full microprogram cycle time was adequately controlled, individual timing pulses varied excessively in duration and placement. Difficulties with the delay-line clock led

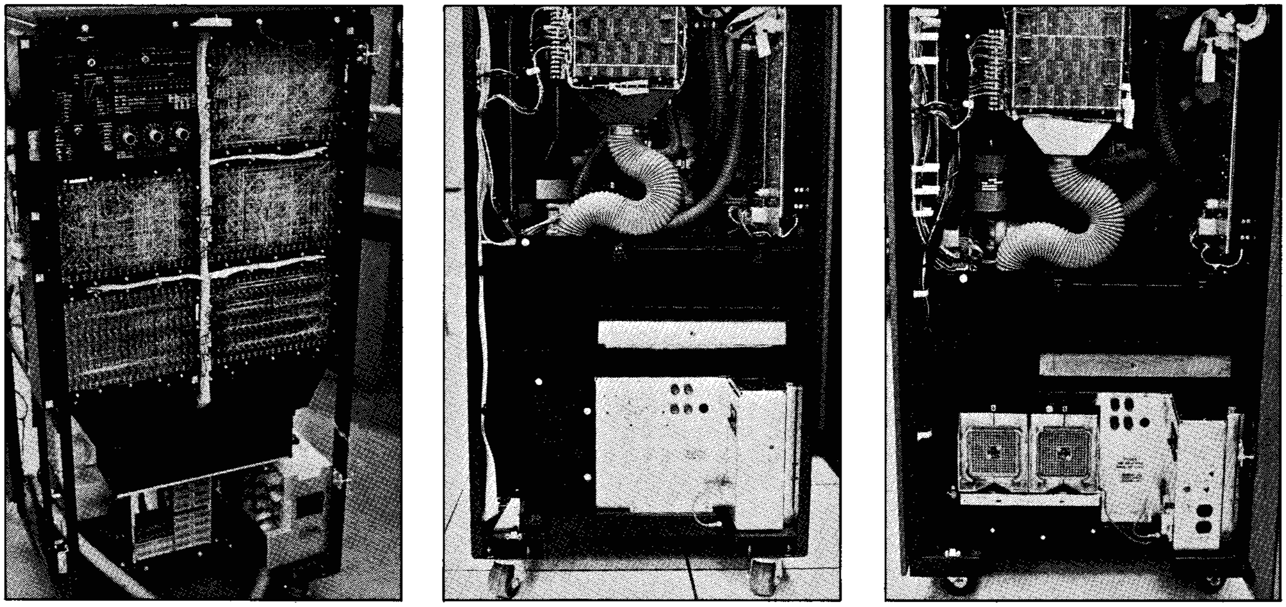

(a) (b) (c)

Figure 3 Tape subsystem packages: (a) Front of IBM 3803 Tape Control Unit, showing CE panel at upper left. (b) Rear of IBM 3420 showing cable connector at bottom left; (c) Rear of IBM 2420 Magnetic Tape Unit, showing cable connector and power supply at bottom (for comparison).

to the adoption of a purely logic clock in later engineering models. The logic clock is of the traveling-wave/phase-hold type, where a pulse is passed from phase-hold latch to phase-hold latch under control of the master oscillator. The clock contains two feedback paths, one of which is gated during each cycle, according to the microinstruction decoded within that cycle. The clock also contains provisions for an immediate suppression of cycles during resets and program traps.

The microinstruction, consisting of 16 data bits plus two parity bits, operates at rates of 150 nsec for the fast instruction to 200 nsec for the slower instruction. The concept is to use an instruction counter and instructions very similar to the instructions used by a processor, i.e., BRANCH-ON CONDITION, BRANCH UNCONDITIONAL, ADD, COMPARE, etc.

The control concept used in the 3803 is to make each microprogram step very simple and rapid so that only one function is performed per step. The instruction counter is preferred to the concept of carrying the next instruction in the current instruction.

Microprogram instructions required for certain optional features to the 3803 are grouped in separate portions of ROS. When the feature is installed these additional portions of ROS are added to the machine along with other hardware required by the feature. A unique microprogram branching instruction permits references to the feature code to be embedded in the basic microprograms but to remain inactive until the feature is physically installed.

- *Physical package*

It would not seem possible to fit the 3803 into a frame the size of a tape drive. In addition to the logic gate and power supplies for the control unit, there is an ac power distribution control for eight tape drives. Thirty-five cables must enter the bottom of the frame: eight for the channel interface, eight for the tape drive interface, eight for tape drive power, eight for switching interconnections, two for power control and one 60-ampere main power cable.

The size of the logic gate presented a major problem. The standard gate, although narrower than the 30-in. over-all width of the control unit, is wider than the frame. This gate is mounted so that it can project past the frame on each side into the hollow portion of the machine cover, as shown in Fig. 3. This approach utilizes space wasted in earlier units.

The size of the power supply presented a second major problem. The past approach was to build the supply as one "box" of generally rectangular dimensions and to size the over-all frame accordingly. In the 3803 the power system is broken into several irregular shapes and tailored to the available space around the cable entry. A notable feature of the power system package is the mounting of large capacitors in blocks of foam plastic to simplify construction and reduce the cost of the unit.

Device interface

The new device interface which connects the tape units to the tape control unit has many advantages over the one

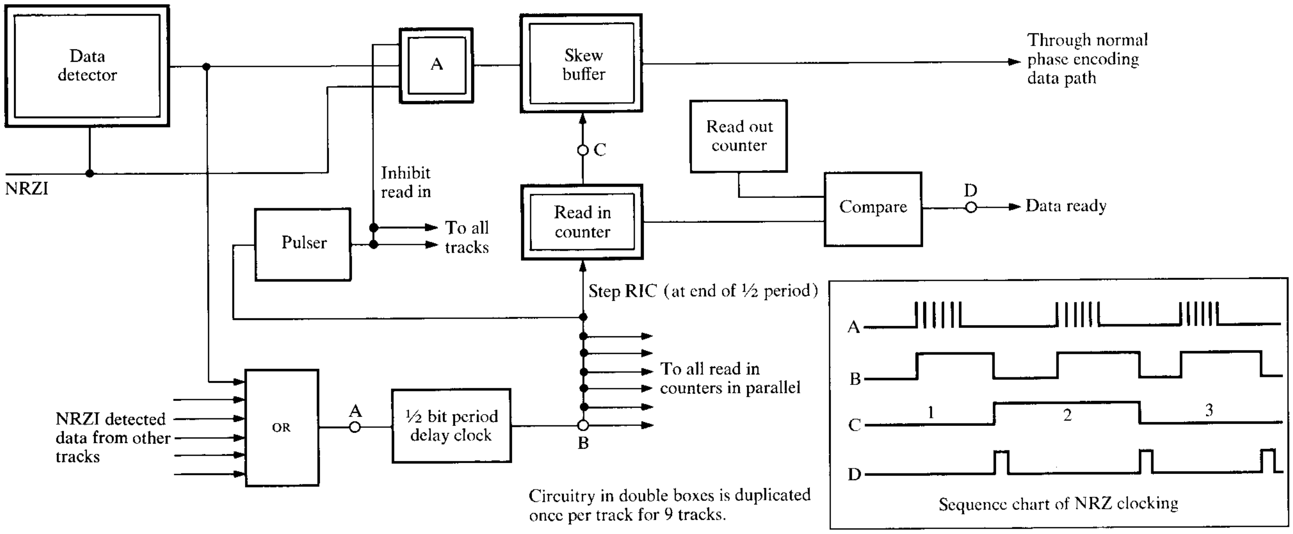

Figure 4 NRZ detection with conversion of NRZ to pseudo-phase encoding.

used in previous subsystems. The 3803/3420 data interface is digital, whereas the 2803/2420 used an analog interface. The significant difference is in the fact that the higher voltage of the digital interface gives better noise rejection characteristics.

The radial connection used on the new subsystem, indicated in Fig. 1, permits a tape drive to be taken offline without disturbing the operation of the other tape units on the subsystem. These tape units are not disturbed nor need they cease operation while a malfunctioning drive is removed or a repaired drive is reconnected.

Subsystem features

The operational features described here have been modified to provide functions similar to those of the predecessor, the IBM 2803/2420/2816, in a far more compact physical configuration and with substantially improved price performance. The modifications indicated in the following paragraphs involved a redesign of the dual-density recording method, the channel and tape unit switching arrangements, and a broad extension of the diagnostic routines.

• *Dual-density feature*

The dual-density feature allows the system to read and write both phase encoding and nine-track NRZI recording. One purpose of this feature is to permit the use of existing nine-track NRZI tape libraries. It also allows the user to interface with systems that do not have phase encoding.

The main difference between the present and previous configurations of the dual-density feature is a drastic simplification of the data channel. The IBM 2415 and 2803 used a dual path consisting of separate data channels for phase encoding and NRZI. In the current design of the 3803, the NRZI data are converted to resemble phase encoded data so that the path for phase encoding can be used for both forms of recording. The conversion (Fig. 4) was accomplished by developing a clock from the NRZI data that is used to drive all the phase-encoded track clocks in synchronization. The elimination of the parallel data path that was used in the 2803 not only deleted several data registers but, more importantly, reduced the need for interconnection between printed circuit cards.

• *Seven-track and nine-track operation*

The IBM 3420/3803 subsystem has a seven-track capability in addition to the basic nine-track phase encoding mode. The addition of the seven-track mode permits tape communication with older systems, such as the IBM 7000 or 1400 series. The 3803 operates in both the seven- and nine-track modes. The 3420 Magnetic Tape Unit is equipped for either seven- or nine-track operation, according to the recording format specified.

• *Channel-switch and tape-unit-switch operations*

The IBM 2803 Model I made two-channel switching available only to the NRZI user. The 3803 now makes it available also to phase-encoded users, improving the flexibility of the computer by giving alternate retry in case of READ error and a way of transferring information from one system to another. Attachment can be made either to two channels on one processor or to single channels on two different processors.

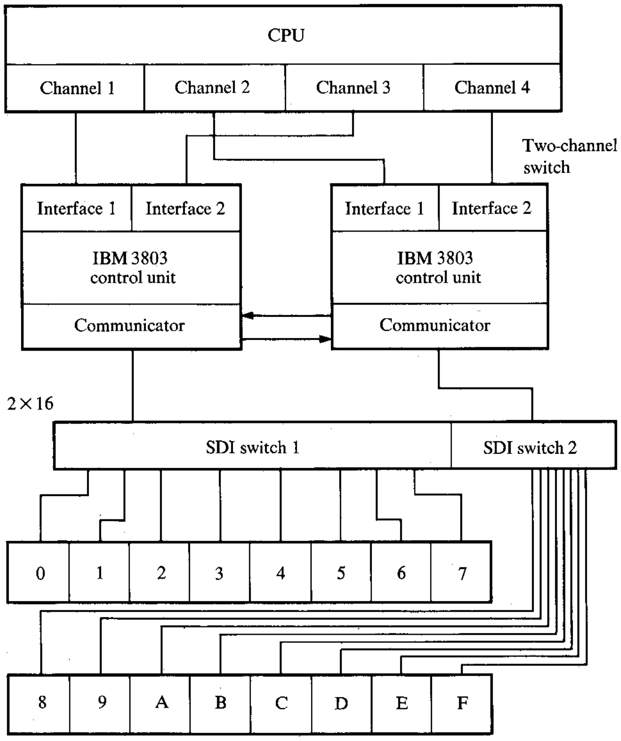

Figure 5 Typical two-channel switch system including 16 tape drives attached.

The new configuration also features a device interface that is switchable. In the basic configuration, one 3803 controls up to eight tape units. Either two, three or four of the basic 3803's can be interconnected for control of 16 tape units in the 2 × 16, 3 × 16, or 4 × 16 configurations. In addition, any or all of the control units can be connected to multiple channels utilizing the two-channel switch. A 4 × 16 scheme is shown in Fig. 5.

• *On-line diagnostics*

The 3803 has greatly expanded the on-line diagnostic testing capability over that of the 2803 and other previous tape control units. The general function is shown in Fig. 6.

The improvements are the result of implementing diagnostic programs and checks into the microprogram. Including these abilities in the microprogram relieves the CPU from the burdens of timings and gives diagnostic capabilities that are transparent to the CPU.

Some examples of these routines are given below:

a) *Internal timing loops.* One important characteristic of previous diagnostic packages was to measure the interblock gap. This was accomplished by counting the number of CPU instructions between blocks to give an estimate of distance. The 3803 has a routine in the microprogram that counts the tachometer pulse and gives an accurate measurement of distance.

b) *Additional sense bytes.* Older tape control units have only six sense bytes to convey subsystem status back to the CPU. The 3803 has expanded this function to 24 sense bytes with greatly expanded capabilities. An example of this is the ability to sense a tape unit which has gone "not ready" and indicates the cause.

c) *Microprogram checking.* During normal operation the microprogram monitors tape unit characteristics such as start time and steady state velocity. During diagnostic operation additional checks are performed and expanded to give a profile of start/stop times and the thread and load sequence. A quality check can also be performed on the tachometer pulses.

d) *Loop* WRITE *to* READ. A serviceability restriction on previous subsystems was the inability to separate malfunctions in the tape unit mechanics from malfunctions in the READ/WRITE paths. In order to overcome this the 3803 has built into it a loop WRITE-to-READ feature which checks the READ/WRITE data paths, including the switch out to the tape unit and back. This operation can be performed on-line with a diagnostic program.

A sense-byte analysis program is available which reduces all the accumulated sense data to meaningful information so that the field engineer can in many cases go directly to the failing card.

The IBM 3420 Magnetic Tape Unit

Although the outward appearance of the IBM 3420 is very similar to that of its predecessor, the 2420, there are several new design concepts, which will be described here. Several constructional features have been carried over from the 2420: the low-inertia motor for capstan drive, automatic cartridge loading, and soft tape handling, in which only the READ-WRITE head touches the oxide side of the tape. The new design modifications include a far more sensitive and more reliable method of phase-error detection and correction; the use of microprogramming for ROS control; a more complete READ check; a new monolithic READ/WRITE detection channel; substantially better control of reel speed and improved rewind performance through the use of additional tachometers and fiber-optics for tape speed sensing; and an automatic latch for seating and locking the tape reel. Operating characteristics for three models are indicated in Table 1.

Some design innovations are described in the following paragraphs.

• *Phase-error detection and correction*

In previous tape subsystems the error detection for phase-encoded recording has been primarily based on comparison of the signal amplitude being read from the tape to a predetermined reference level. Any signal falling below this level has been the identifier of an error. In

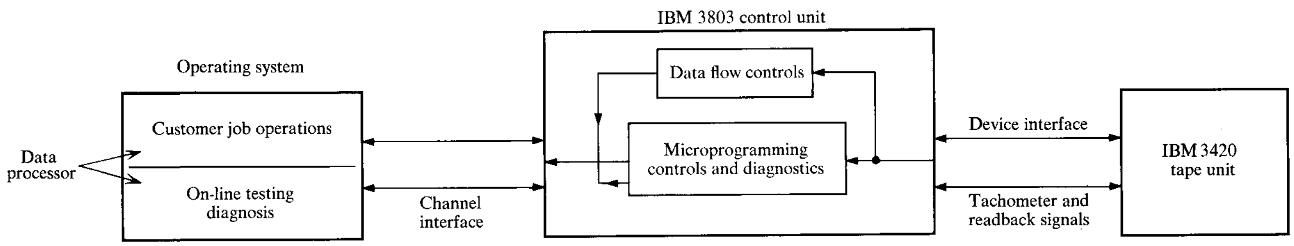

Figure 6 The on-line diagnostic function of the IBM 3803/3420 Magnetic Tape Subsystem.

Table 1 Some operating characteristics of Models 3, 5 and 7 of the IBM 3420 Tape Unit.

Model Number	*Data rate (Kilobytes/sec)* *1600 bpi*	*800 bpi*	*556 bpi*	*Tape speed, (in./sec)*	*Nominal IBG time, (msec)* *7-track*	*9-track*	*Nominal rewind time, in sec for 2400-ft reel*	*Nominal read access time (msec)*
3	120	60	41.7	75	10.0	8.0	70	4.0
5	200	100	69.5	125	6.0	4.8	60	2.9
7	320	160	111.2	200	3.75	3.0	45	2.0

the present system no such level is required, since its detection of an error is based on the phase properties. This method permits the reading of signals that have degraded to an extremely low amplitude level. Figure 7(a) represents the permissible phase distortion within the specified phase gate. Figure 7(b) shows the phase gate circuitry used in a single chip module. Regardless of the amplitude of the signal, and as long as the phase shift has not exceeded the gate boundaries, an error will not be flagged and the data will be read accurately.

The boundaries of the phase gate for reading are generally set between 70% and 130% of a normal bit cell timing. The exact choice of these limits depended on a compromise among the characteristics of the tape recording surface, head and read circuitry. The same circuitry is utilized during the WRITE operation. However, the phase gate bounds for writing are reduced to 70% and 125%, thus ensuring that the data have been written with less phase shift than will be allowed on readback.

An additional constraint is placed on the writing process in the form of an amplitude sensing threshold so that the recorded signal meets industry standards.

Phase-error detection is performed on a per-track basis and points to the track in error for error correction. The over-all performance of error recovery is enhanced considerably because the phase-error method permits reading data with signal amplitudes approaching 5% of the nominal signal.

- *Full readback parity check*

In previous phase encoded tape subsystems readback check consisted chiefly of an amplitude check of the re-

Figure 7 (a) Phase gate generation. (b) Phase gate circuits.

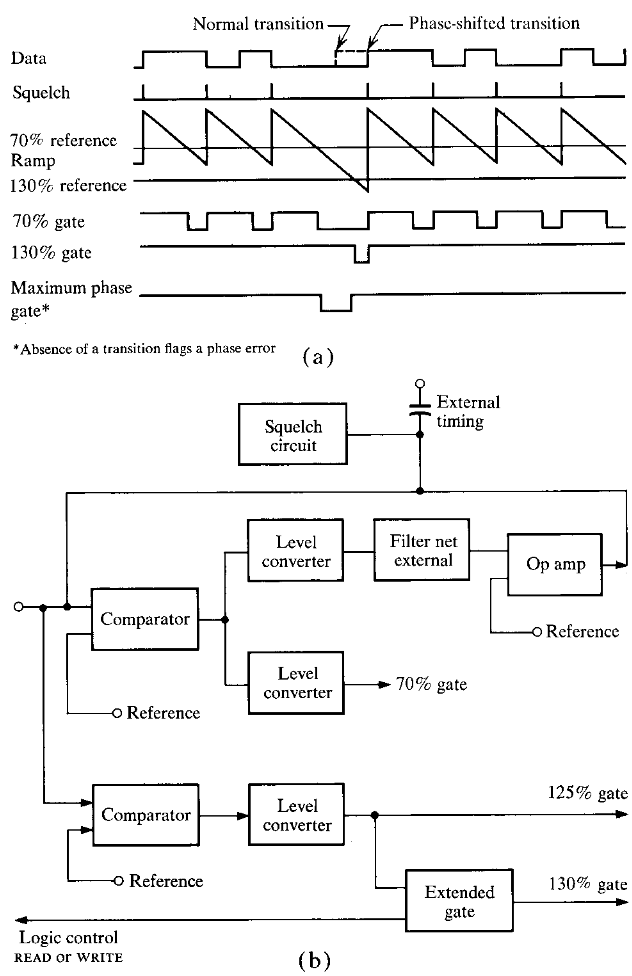

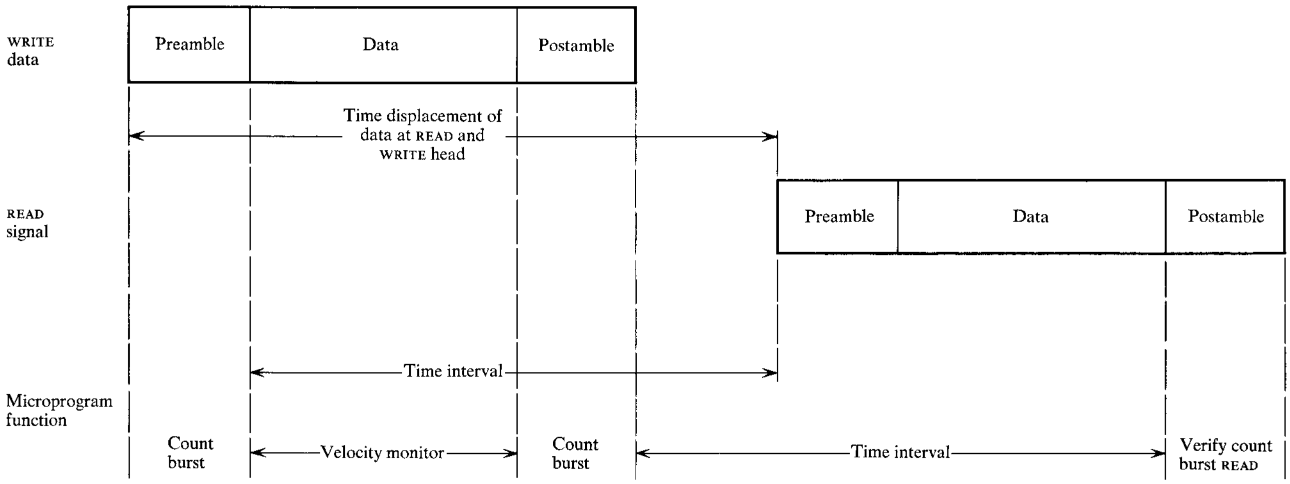

Figure 8 Microprogramming monitor of WRITE data format.

corded data. In the present design the error checking is expanded to include data detectability (phase errors and vertical redundancy) and proper formatting of the synchronizing burst of markers as well as the conventional amplitude check. The formatting check is executed by a combination of microprogramming and clocking circuits.

The monitoring of the WRITE data format by the microprogram is illustrated in Fig. 8. The control and counting functions for writing the phase encoding format bursts are provided by the microprogram and are executed by the write hardware. During the data writing portion of the block, the write hardware operates independently, which frees the microprogram to monitor tape drive velocity. The microprogram also makes a series of timing checks that measure the relation between the beginning and end of the WRITE operation on tape to the beginning and end of the data as seen at the read head. Any significant deviation is then signaled as an error. In addition the microprogram makes an actual count of the number of bytes in the trailing burst to ensure that the data can be read reliably.

These microprogramming steps not only provide a fuller and more dependable READ check but also prevent catastrophic runaways that are due to malfunctions in the write hardware.

• *Use of incremental encoders in reel control*

In order to obtain more accurate information on reel speed than was available in the IBM 2420, a tachometer was added on the reel side of each vacuum column. In order to minimize inertia and to retain high reliability, an optical incremental encoder was adopted to sense tape speed at the reel. It was necessary to sense actual tape speed rather than reel motor speed because of the large changes in reel radius during operation.

The high-resolution reel speed information is used in two ways: During normal-speed operation the tachometer is used to prevent overspeed of the reel motor and limit worst-case excursions of the tape loop in the column. In rewind, the tachometers are used to help control capstan speed, allowing the capstan to accelerate to the speed the reels are capable of following.

The radius and line density of the tachometer idler were chosen to give a frequency equal to 1/16 of that of the capstan tachometer when the tape speed at the reel is equal to capstan speed. This technique allowed use of a pulse-rate subtractor to determine the loop speed and direction.

Normal speed operation. Figure 9(a) shows the velocity of the tape at the capstan and reel during a worst-case reversal of the capstan. "Worst case" is defined as that sequence of capstan operations that produces maximum excursion of the tape within the column. After the capstan has been moving tape at +200 in. for about 100 msec, the reel has overshot capstan speed, thus starting the tape loop back toward the center of the column. The tachometer is designed to sense the reel overspeed before the reel reaches 114% of capstan speed. Worst-case excursion occurs if the capstan is stopped at the moment reel overspeed is sensed. A low-torque drive of the reel motor accelerates the tape to a maximum of 133% of normal capstan speed and just as the tape loop enters the center brake zone of the column, the capstan is reversed and is driven to −200 in./sec. The tape loop moves toward the other end of the column, but keeps a generous distance from the end in the worst case because of the reduction in reel overspeed made possible by the reel tachometers.

Digital simulation studies at this laboratory have shown that the use of the tachometers provided an addi-

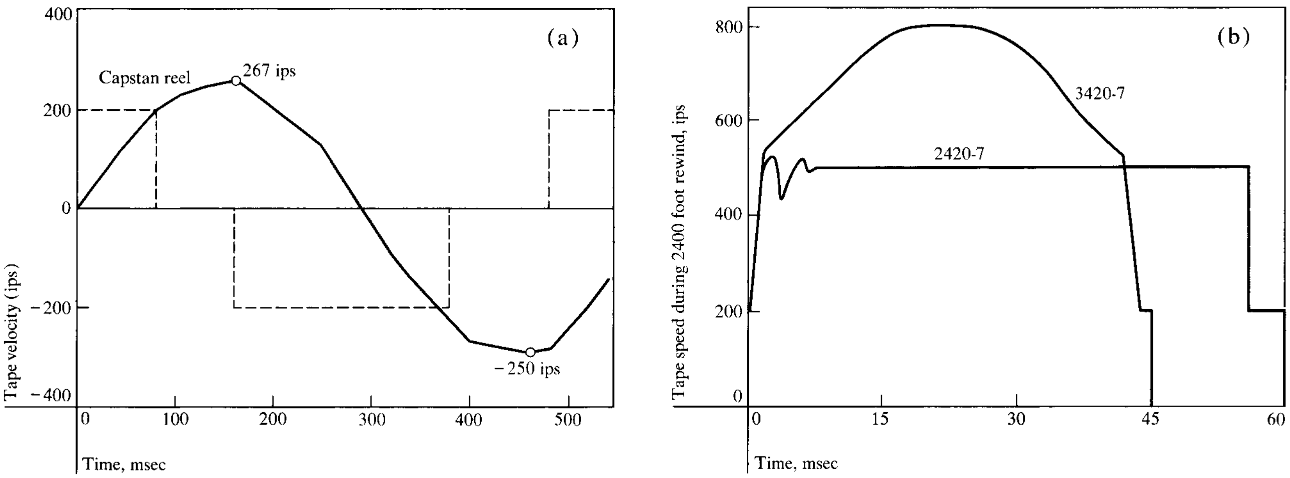

Figure 9 Tape velocity data. (a) Reel speed during worst-case capstan reversals. (b) Rewind speed characteristics.

tional six-inch reduction in maximum loop excursion. This was later verified in machine tests.

Rewinding. The rewind performance was also improved by use of the reel tachometers. Previously the entire rewind was done at some controlled speed lower than the minimum speed at which the empty reel can be driven, as shown in Fig. 9(b). In the 3420, on the other hand, rewind speed is not fixed. Instead, the capstan is driven to whatever speed the weaker reel motor is capable of reaching. The reel tachometers are used in conjunction with loop position sensing switches to determine which of the reel motors is the weaker and to keep the speed of the two reel motors and the capstan motor together. A maximum speed is reached [as shown in Fig. 9(b)] when the radius of tape on each reel is about equal. Another improvement offered by the tachometers is the gentler deceleration of the reels when the end of rewind is approaching.

Fiber optics. The single light source for photosensing in the 3420 is placed in a location remote from the tape path and other temperature-sensitive areas. The light is piped by means of fiber optic bundles to the areas where reflective markers on the tape are detected by photodiodes to determine tape position. Fiber optics are used in the three tachometer feedback networks and in the take-up reel rotational counter circuits.

• READ/WRITE *detection channel configuration*

The 3803/3420 READ channel has been designed and implemented using the monolithic technology. Four monolithic modules were designed for the tape subsystem. These are indicated in the block diagram of data channel diagram, Fig. 10: READ chain, amplitude sense and gating ramp generator, and time sense.

Figure 10 Data channel diagram.

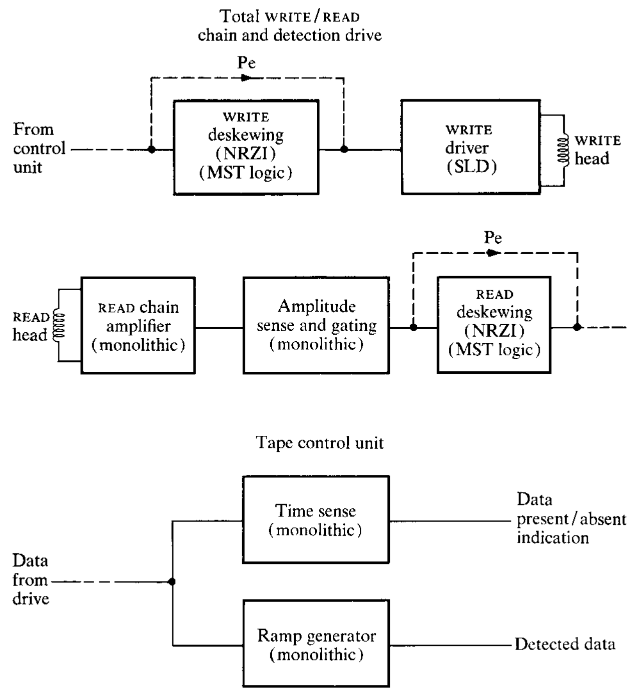

The READ chain module indicated in Fig. 11 is a design in which the function is to amplify, differentiate and limit magnetic tape signals at frequencies of 10 to 600 kHz. It is a direct-coupled, differential amplifier with ac coupling in the differentiating stages having a provision for stage balancing through a differential dc voltage feedback circuit. Frequency response and amplifier gain is controlled externally for flexibility in machine applica-

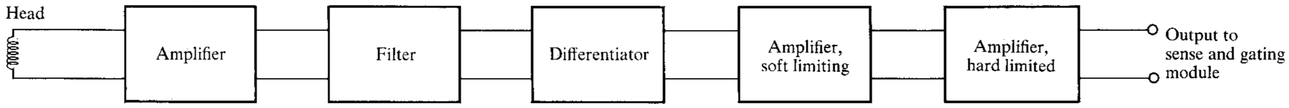

Figure 11 Block diagram of READ chain.

tions. Using a differential design approach provides a considerably higher signal/noise sensitivity.

The significant advantage of this design philosophy over preceding tape subsystems is the reduction in packaging size and power requirements. Both component reliability and signal/noise sensitivity for the differential design were improved over that of previous tape units by a factor of 5.

Conclusions

This paper has indicated how the microprogrammed read-only storage control, the new phase-error detection/correction scheme, the full readback parity check, and the improved reel control represent a combination of significant engineering improvements. In addition, the physical size of the IBM 3803 ranges from 1/2 to 1/4 of the size of the 2803 and at the same time offers an appreciably improved cost performance.

The serviceability of the subsystem has benefited from two aspects. The first is the introduction of the radial connection scheme, which permits continued operation when a tape unit is being removed for repair or maintenance. The second is a new version of on-line diagnostics, in which sense information gives a picture of the subsystem status at any point in time, based on 24 bytes of information instead of the six bytes used in previous models.

The factors contributing to the over-all system performance are the increased reliability of electronic components, the existence of fewer replaceable units and the improved methods of verifying written data.

Acknowledgments

We would like to acknowledge the combined efforts of the design team responsible for the development of the tape subsystem. Special thanks are due to Richard Van Pelt for his work on the reel controls, to the Circuit Technology group for design layout and implementation of the read detection module, and to Richard Breitenbach for his contributions to on-line diagnostics. In addition we acknowledge the support and technical direction of Jack F. Wells. The basic read detection module was developed in the IBM Components Division, East Fishkill laboratory.

Received December 11, 1970

The authors are at the IBM Systems Development Division laboratory, Boulder, Colorado 80301.