DESIGN PRINCIPLES OF THE MAGNETIC TAPE SYSTEM FOR ATLAS COMPUTER

by

D. B. G. Edwards, D. Aspinall and T. Kilburn

Manchester University

Manchester, England

INTRODUCTION

The ATLAS digital computing machine is capable of a computation rate of $5 \cdot 10^5$ operations per second, a high rate which demands an extensive memory. The main store of the ATLAS at Manchester University is provided by 16,384 words of 48 digits on magnetic cores, backed by a magnetic drum containing 98,304 words. These two stores are presented to the programmer as one level of storage of 98,304 words [1]. The storage can be extended on other versions of the ATLAS and augmented by large file drums or disc stores. There is still a need for a substantial bulk store which is provided for by a library of magnetic tapes. These tapes serve different purposes in the machine and in order to carry them out simultaneously there is a need for several parallel channels all of which can operate at the same time. Three channels are used to accomodate the operating system of the machine [2], one to accept all the input data from various paper tape and card equipment, one to hold output information which is being fed to tape punches, parallel printers etc. and the third to act as a "dump" for programmes in abeyance. Different problems require different amounts of computation and input/output time and on any one problem these times cannot be overlapped efficiently. The significance of the above operating procedure is to provide a delay in the system so that these times overlap over several problems.

Certain problems within the machine will have input data in the library of magnetic tapes which must be updated, an operation requiring an extra three channels; one for the old record, one for the updating information and one for the new tape. A further two channels are provided to transfer data between the library of tapes and the other peripheral equipments such as Line Printers or Card Readers. This makes a total of eight channels in all.

When the transfer from one deck is completed and its reel is being rewound the channel can be freed, to be used by another transferring tape deck. The switching between a channel and several decks is accomplished

in a special Telephone Exchange, the cross points being transistor switches. The most usual arrangement is for two channels to be shared among eight tape decks allowing an installation to include up to 32 tape decks.

For economic reasons the Manchester University ATLAS does not include this exchange, and each of its eight channels has one tape deck. The tape deck is based upon the 1″ version of the Ampex TM.2 tape transport *(Fig. 1)*. The tape speed is 120 inches per second and there are separate

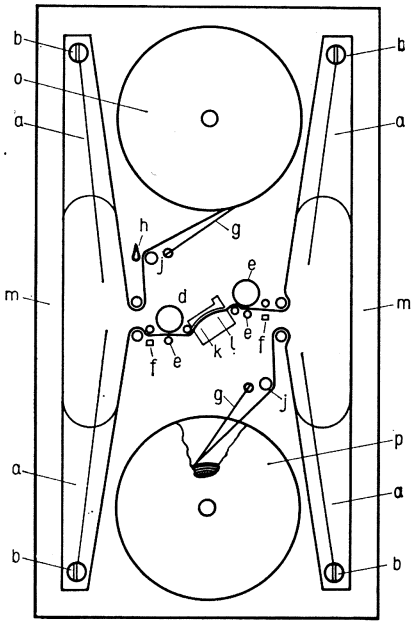

Fig. 1
The tape transport Ampex TM.2.

a — sensing slot, *b* — inlet to vacuum blower, *c* — pinch roller, *d* — capstan for reverse motion, *e* — capstan for forward motion, *f* — tape brake, *g* — pack follower arm, *h* — tape clamp, *j* — sensing post, *k* — write head, *l* — read head, *m* — supply vacuum chamber, *n* — take-up vacuum chamber, *o* — supply reel, *p* — take-up reel

read and write heads each of 16 tracks, a packing density of 375 bits per inch (148 bits/cm) is achieved. The stop time is 1.5 milliseconds and the start time to within 10% of nominal speed is 2 milliseconds. A servo restriction forces a delay of 90 ms to be included between a stop and a direction change command.

The deck is designed to contain the minimum of equipment, as much as possible being shared within the channel. The channels themselves share many circuits within the Tape Coordinator including check sum circuits and a single transfer path between buffer store and the main core store, and a single command path from the central control unit. This is shown in block schematic form in *Fig. 2* which also indicates the inclusion of 1024 words of working store, 8000 words of fixed store, the floating point accumulator

and the peripheral equipments and the magnetic drum which can all operate at the same time as magnetic tape operations.

The many simultaneous operations come under the control of a "Supervisor" [2] routine, most of which is contained in the fixed store. In particular the supervisor routine organizes the search for data on tape and its use in the correct programme. When a programme instruction calls for a particular

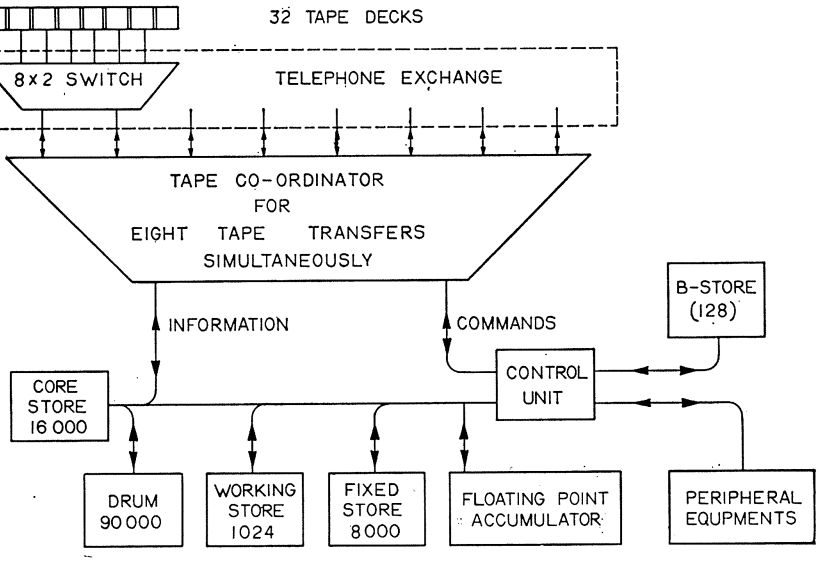

Fig. 2
Block schematic of ATLAS and the tape system

unit of data from tape the supervisor requests the operator to put the reel onto a free deck. There is no need for the operator to select a particular deck or alter the address of a deck; as each new reel of tape is put onto a deck, the supervisor reads the first block of information which contains the identifying factor. The supervisor keeps a directory of the contents of each deck and is able to search for data from tape as requested by the programme. The transfer takes place autonomously including separate parity checks on the transfers between core store and coordinator, and those between coordinator and magnetic tape are checked by a check sum. If a fault is indicated the supervisor organizes a recovery process, for which it is convenient to read in both forward and reverse directions. The write to tape transfer takes place in the forward direction of tape motion so that the

written data can be checked by a trailing read head. Data is written to tape in blocks of fixed length. Each block has an address which is written during an initial testing and addressing operation.

FIXED OR VARIABLE BLOCK LENGTH ON TAPE

The aim of the design has been to give the programmer tape reliability throughout a tape's life, without it being guaranteed dropout free. The programmer prefers no restriction on the length of block on tape but finds certain limitations in practice. If the block is very short then the efficiency of tape use is poor, blocks less than 128 words are shorter than the interblock gap. The upper limit to block length is defined by the size of the main computer store, this could be 98,000 words on ATLAS, but the tape is to be the means of communicating with smaller satellite computers with small main stores which cannot cope with blocks longer than 512 words.

If a block is written to a region of tape which is found to contain a permanent fault then the block must be rewritten further along the tape wasting tape and time. The waste is proportional to the block length and it is economical to divide a large batch of data into blocks.

If variable block lengths are allowed then the same region of bad tape will be repeatedly encountered leading to recurring waste. A fixed block length permits the tape to be pretested with its good regions marked and addressed. Tape which develops a fault can be taken out of service by erasing its marks and readdressing the remaining good tape. In this way the quality of tape used in transfers does not deteriorate with age and time wasted in rectifying permanent tape faults does not increase.

The fixed block system makes it possible to rewrite new information over old without the risk of destroying good information in neighbouring blocks, a facility leading to the use of magnetic tape as an efficient backing store which is exploited for the internal organization of the computer by the Supervisor routine. The main core store and drum both operate with fixed blocks of 512 words and this block length is selected for magnetic tape. Programmers are offered the variable block length facility through subroutines in the fixed store. The variable length blocks are assembled into fixed lengths by the Supervisor routine.

When very short blocks are encountered these are first assembled and remain in a section of the core store until a full block of 512 words is available for transfer. This is slightly wasteful of core store space but is not thought to be significant.

Accurate location of data on tape is guaranteed by the system of block marks and permanent addresses. The address is checked as each transfer takes place. In particular the address is checked before write current is switched on in a write transfer to ensure that valuable data is not overwritten. The accurate transfer of data between the computer and magnetic tape is confirmed by means of a check sum, which is written to tape as an additional half word at the end of the 512 word block.

LAYOUT ON TAPE

The block layout is shown in *Fig. 3*. The separate read and write heads of the tape deck each have 16 tracks allocated as follows:

12 Information
2 Clock
1 Block Marker
1 Reference Marker.

The information within the block is written as a stripe of 12 digits, or two 6-digit characters. Each character has a centrally located clock digit. The block addresses are written in the same way as information but distinguished by means of the mark on the separate Block Marker track. The mark consists of 13 reversals of magnetisation, 3 of which must be detected by an integrator circuit, before the mark is indicated to reduce the possibility of spurious signals on this track being mistaken for the mark. The marks themselves can be used for accurate location of information but it is necessary to have the additional check made possible by the address of each block.

The marks are protected from accidental erasure by switching out the relevant write amplifiers. It is not possible to protect the address in this way. In normal use the write current is switched on only when the tape is moving in the forward direction and the address has been checked. This accounts for a gap *A* of *Fig. 3*. It is made up of distance between read and write heads plus time to check the address plus time for which write current is turned on to produce a margin which erases previous data. The block of 2048 information stripes plus 2 check sum stripes is then written to tape. A signal from the trailing block mark turns the write current off to ensure that the next address is not erased.

Thus gap *B* of *Fig. 3* between trailing block mark of one block and leading block mark of the next must be greater than 0.39 inches, the distance between read and write heads. This gap length also depends upon the stop-

start characteristics of the tape transport, which receives a stop command as a selected block mark is read. The tape is brought to rest so that the read head is in the interblock gap and in such a position that on restart, in either direction, the tape is up to speed when a block mark is read. The gap B must not be smaller than twice the start distance. The start distance is greater than the stop distance and the stop command is given to the deck 1.5 ms after the block mark has been sensed so that the tape may be restarted in the opposite direction.

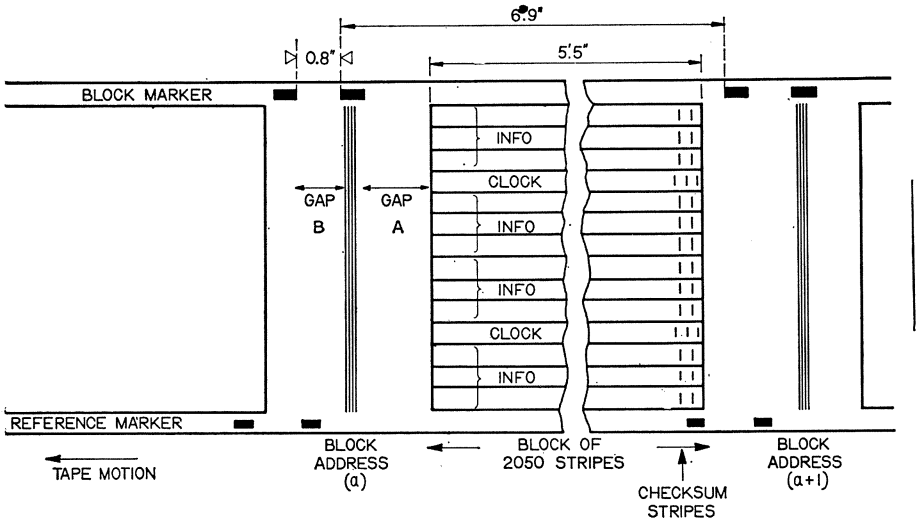

Fig. 3
Layout of magnetic tape

The manufacturers state that the maximum distance covered 2 ms after a start command is 0.18 inches. The tape has a tendency to flap away from the head at this distance unless the pinch roller and actuator assembly is carefully adjusted. It has been necessary to assume a safe start time of 3.5 ms during which time the maximum tape movement is 0.380 inches. This allows the tape transport adjustment to be less critical and reduces the amount of preventive maintenance.

The block marks exist not only to distinguish between address and information but also to stop the tape in the inter block gap and define the regions of proved and tested tape which is available for information.

The block address is aligned in four stripes adjacent to the leading block mark. The mark and address are written during the initial addressing

process and rewritten if the tape is re-addressed. The location of the correct region of tape to write a block mark is decided by the reference mark. During the initial addressing process this mark is written to tape and when it is read the leading write head puts the block mark to tape. The positioning of consecutive reference marks to define the inter block and block lengths is decided by a crystal controlled counter in the tape coordinator. In the re-addressing process this mark is read and the write head puts the block mark to tape. If a block is faulty its reference marks are erased and no block marks can be written in the bad region when the tape is re-addressed.

The addressing and re-addressing processes are carried out on one of the channels of the tape coordinator under programme control. This operation can proceed at the same time as transfers on the other seven channels. Each installation has its own mechanism for producing addressed tapes, which can operate when the computer tape system is not busy, making it independent of supplies from a distant source.

The tape is tested during the addressing process by writing a pattern of ones to tape in the region of the block. This pattern is checked by the trailing read head and a single fault condemns the block.

WRITE ELECTRONICS

Write current is turned on under programme control after the block address has been checked. Even so, valuable information may be over-written under fault conditions. To prevent this a "write permit" ring is supplied which must be added to the reel before its attachment to the deck to allow write current to occur. A reel which has just been written may require immediate protection. A "write inhibit" button is provided on the deck which allows the operator to immobilize the write circuits by switching off the write power, even though the "write permit" ring is present. Other logical interlocks are provided to prevent the write command being obeyed when tape is moving in the reverse direction or at fast rewind speed. Furthermore every block mark signal resets the write command.

The write head has the following dimensions:

gap width = 0.0005 inches (12.7 μm)
track width = 0.031 inches (0.79 mm).

The write current needed to saturate the tape is 80 mA, but a write current of 120 mA is used to ensure adequate erasure of old information.

The digit representation on tape is modified non return to zero in which each flux reversal represents the digit 1. This system is simple to use since

a reversal of either polarity read from tape represents a 1. The write electronics must include a counter which changes state as it receives a pulse representing 1. The output of the counter governs the direction of current in the write head and is in the tape deck receiving signals from the distant tape coordinator as shown in *Fig. 4.*

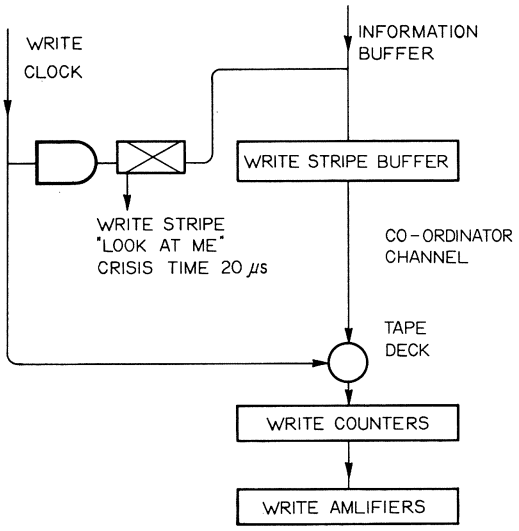

Fig. 4
Write electronics

Each channel has a Write Stripe buffer register which holds 12 digits. It is written to from a word buffer store within the Tape Coordinator at a time when its output is not being gated. The Write Clock interrogates the output of the write stripe buffer every 22.2 μs, under the control of a crystal oscillator. This gating takes place in the tape deck so that all 12 write counters set at the same time, independent of differing cable delays between the coordinator and the deck. The write clock pulse is delayed by 2 μs to allow the write counters to set, before setting the Write Stripe "Look at Me" flip flop. This flip flop indicates to the Tape Coordinator control circuits that this channel has just written a stripe to tape and demands the transfer of a new stripe into the write stripe buffer within 20 μs. This time is termed the "Crisis Time" of the write stripe buffer.

READING FROM TAPE

Each read head has the following dimensions:

gap width = 0.00025 inches (6.3 μm)
track width = 0.024 to 0.025 inches (0.61—0.63 mm).

It gives an output signal of approximately 15 mV peak to peak from 3 M's 499 magnetic tape. The amplitude will vary due to poor tracking, dust particles between the head and the tape causing the tape to lift off the head

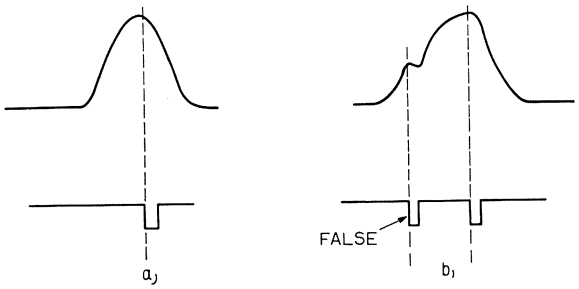

Fig. 5
Peak detection waveforms: *a)* noise free, *b)* noise present

and the tape flapping away from the head immediately after starting to move. The read signal is detected by a method which is independent of absolute amplitude, such as Peak detection *(Fig. 5a)*.

On installations which include a Telephone Exchange it is economic to put much equipment in the Tape Coordinator which can be shared between several decks. For this reason the detection circuits are in the Tape Coordinator and the Tape Deck contains simple circuits to amplify the read signals to a level which can be transmitted through a cable and telephone exchange.

Noise due to interference which is picked up by this cable causes distortion of the read signal. The high frequency noise superimposed upon the read signal can cause shift in time of the peak and false signals as shown in *Fig. 5b*. The peak detection circuit gives an output when the input voltage has zero rate of change. Noise can cancel the rate of change of the input signal and cause zero rate to occur and be detected away from the actual peak if the detector has a high frequency response.

Thus the high frequency noise must be removed by means of a low pass filter. This measure seems to be adequate on the Atlas Magnetic Tape

System where the maximum signal frequency is 45 Kc/s. However the Magnetic Drum system has a read signal of higher frequency (500 Kc) closer to that of the noise (4 Mc/s) and a different detection scheme is used.

The principle of this angle detection scheme is illustrated in *Fig.6a.* A delayed version of the signal is subtracted from the signal itself and the point of zero difference is detected and used to indicate the presence of the

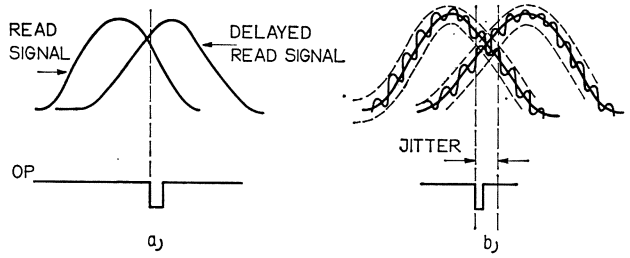

Fig. 6
Angle detection waveforms: *a)* noise free, *b)* noise present

signal. The detection point is a small angle away from the peak. It must be near the peak to reduce shift of the detection point due to waveform assymetry. Noise causes shift of the peak but has much less effect upon the point of zero difference as illustrated in *Fig. 6b* and described in Appendix *A*.

TIME DISPLACEMENT OF READ SIGNALS

The minimum time interval between reading stripes from tape is important in the design of the system and its calculation follows. When the tape is written its speed may vary $\pm 12\%$ from the nominal 120 in/s, shortly after start. The important case to consider for the purpose of this calculation is a slow speed on write of 105.6 in/s and a fast read at 134.4 in/s when the digits will be read at a nominal time interval of 17.4 microseconds. This time can be reduced by a combined jitter of write electronics and peak detection circuits accounting for a jitter of ± 1.4 microseconds in the read signal. Thus the time between read signals on the same track can be as short as 14.6 microseconds.

The output signal from a peak detection circuit may be displaced in time from other signals of the same stripe due to jitter and also tape skew and variations in the mechanical construction of the recording heads.This displacement is greatest when the tape is written fast and read at slow speed on different decks with opposite variations [3].

The time displacement of digits in the extreme tracks of a 6-digit character in the worst case is ± 5.95 microseconds, about a mean time corresponding to an ideal clock signal written at the same time as the stripe. (This time for the slow write/fast read extreme is ± 5.05 microseconds.) Each read signal is stored on a flip flop of the 'A' register *(Fig. 7)* until all digits of the character have been read then all 6 digits are transferred

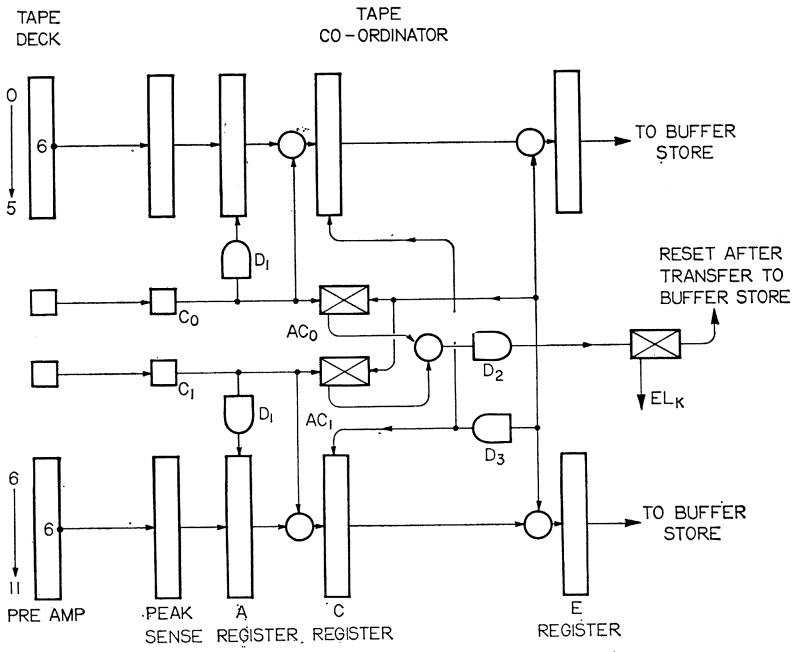

Fig. 7
Read electronics

simultaneously to a second register (the *C* register). It has been customary to use a delayed clock signal to strobe this transfer. If the clock signal is ideal the delay must be greater than 5.95 microseconds and less than $[17.4 - 5.95] = 12.35$ microseconds to account for both extremes of tape speeds. The clock signal may be displaced in practice by ± 3.05 μs in the case of fast write/slow read, and 2.9 μs for the other extreme. Thus the delay must be greater than $5.95 + 3.05 = 9.0$ μsecs and less than $(12.35 - - 2.9) = 9.45$ microseconds requiring a delay circuit of 2% accuracy. An alternative solution is to write the clock digit half way between stripes controlled by a crystal oscillator as shown in *Fig. 8a*. This allows the direct

output from the clock track peak detection circuit to be used as the strobe. The timing of this signal with respect to the digits of the character is shown in *Fig. 8b*. It is seen that the earliest clock signal is later than the latest information signal and the latest clock signal is earlier than the information signals of the next character, by 0.75 microseconds. This time increases for slower read speeds or faster write speeds.

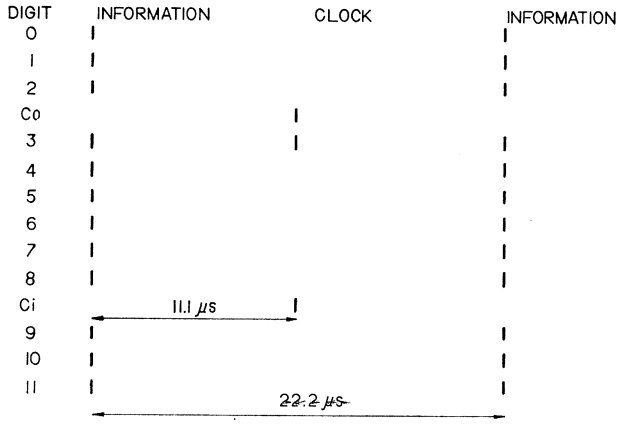

Fig. 8a
Timing of write signals

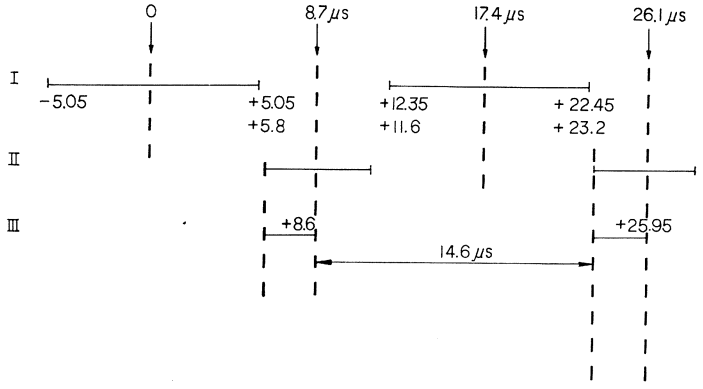

Fig. 8b
Time displacement of read signals
I — displacement of read information on digit tracks 0—5
II — displacement of read clock C_0
III — jitter on clock C_0.
Write speed = 105.6 in/s, read speed = 134.4 in/s

The jitter on the clock track causes the character to be transferred to the C register at intervals as short as 14.6 microseconds in which time the character must be cleared. There are 16 such registers in the Tape Coordinator and it is not possible to transfer from all of them in 14.6 μs without additional buffer storage. It is necessary to have a third register, the E register to which the stripe is transferred after the later of the two characters has been set in the C register as shown in *Fig. 7*.

The clock signal C_0 sets a flip flop AC_0 and transfers the character from the A to the C register. It is then delayed by D_1 to reset half the A register. This is performed by standard ATLAS circuits and takes approximately 0.2 microseconds, much less than the 0.75 μs allowed. The output AC_0 is gated with a similar signal AC_1. The later clock pulse produces a signal from this gate which is delayed by D_2, to allow the C register to set and then transfers the twelve digits of C to E also resetting AC_0 and AC_1. After a further delay D_3 the C register is reset.

The flip flop EL_k is set as the E register is set. It indicates to the control circuits of the Tape Coordinator that a stripe is ready for transfer to the buffer store, furthermore the transfer must take place within 14.6 μs the crisis time before the next stripe is read into the E register.

THE BUFFER STORE

The buffer store consists of 16 words of 50 digits, (48 digits for the information and two for parity) on magnetic cores [4] with a cycle time of 1 microsecond. Each channel has two words of storage and the stripes are packed as shown in *Fig. 9a*. On reading from tape a word is filled from the stripe register and when full, a word "Look at Me" is set. This word waits to be transferred to the main core store whilst stripes are transferring into the second word of storage. Each word has its own "Look at Me" so that for a short time two words can be waiting, as in the example shown in *Fig. 9b*, when a tape of high packing density is being read at highest speed. The last stripe of word Wx is set into the E register at time d, assuming the tape coordinator control takes the longest permissible time to transfer this stripe it is not set into the buffer store and the word "Look at Me" Wx set until time D, 14.6 microseconds later. The shortest time between d and setting the first stripe of Wx into the E register at a is $4 \times 17.4 + 14.6 =$ $= 84.2$ μs as shown. The stripe may be transferred to the buffer immediately at A by which time the word must have been cleared to the main core store. The maximum permissible time between forming the word in the buffer and transferring to the main core store is $84.2 - 14.6 = 69.6$ μs. This time

is called the Read Word Crisis time. The Write word crisis time is much longer, being $(5 \times 22.2 - T)$ when T is the longest time to deal with a write stripe crisis. If $T = 20$ μs then the write word crisis time is 91 μs.

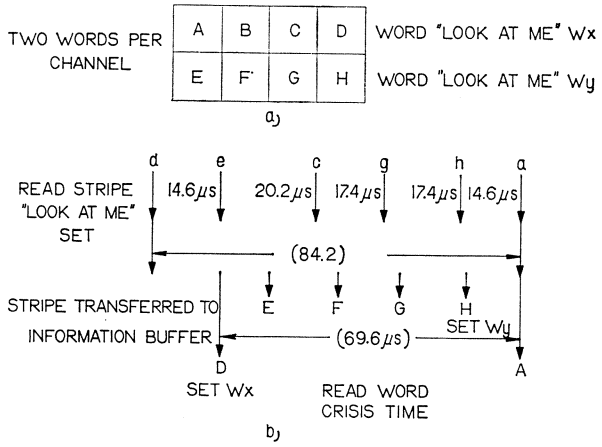

Fig. 9
Read word crisis time: *a)* words in information buffer, *b)* timing of read transfers: worst case

METHOD OF TIME DIVISION

The buffer core store is shared between eight channels. Its operations consist of stripe and word transfers which need not occur at the same time. Other operations which take place during a stripe transfer are; check on parity and formation of check sum on writing to tape, and checking check sum and formation of parity on reading from tape. These use circuits which are shared between the eight channels.

The method of sharing which is used is to divide time into discrete intervals each allocated to a particular channel. If the channel requires attention during this interval it is dealt with. If not, then the time is wasted. The duration of this time interval t can be estimated as follows: The worst case to consider is that of read transfers on all eight channels

Word crisis time $T_W = 69.6$ μs
Stripe crisis time $T_S = 14.6$ μs

$$\text{then } \frac{8t}{T_w} + \frac{8t}{T_s} < 1, \quad t < 1.5\ \mu s.$$

This does not take into account the wasted time, which reduces t to 1.2 μs.

The core store has a cycle time of 1 μs, the adder circuit [5] used in the formation of the check sum has an operation time of 0.2 μs and the parity circuit also operates in 0.2 μs. There is time enough for the operations to be completed within the time allowed and the synchronous system is relatively simple to commission and maintain.

Under fault conditions transfers can be missed. Stripe transfers which do not take place will be detected by the check sum system, and missing words are discovered by a counting circuit. There is also a system to check that word transfers between coordinator and main core store do take place as requested, faulty main store transfers are detected by the parity circuits.

CHECK SUM

The check sum consists of 24 binary digits. It is formed during the write transfer by pairing consecutive stripes into 24 digit half words and adding them all up with an adder including end around carry. This sum is written to tape as two stripes after the block and also stored on a register in the tape coordinator. During the write transfer the trailing read head detects the stripes which are totalled to form a read check sum. The check sum achieved in this way is compared with the last two stripes read from tape and the stored check sum. If all three agree then the transfer is successful.

CONCLUSIONS

The tape system will provide an efficient form of bulk storage and input/output medium for the ATLAS computer. Its preaddressed tapes of fixed block length will lead to little loss of computing time due to permanent tape faults. The time division method is simple to maintain. This has been achieved by working well within the limits of the circuit technique.

Though it is possible to achieve improved performance from the magnetic tape medium by using advanced circuit techniques, this could not be matched by the relatively poor reliability of the tape transport. It is important to improve the reliability of tape transports before the advanced performance can be exploited in a computer system.

At present the computer is addressing its own tape and data about tape wear is being acquired. The complete system of 8 tape decks will be in operation by December 1962, when performance figures will be obtained.

ACKNOWLEDGEMENTS

The authors wish to thank the Atlas Computer Group at Manchester University and Ferranti Ltd., for their assistance at all stages of this project.

APPENDIX

Angle Detection

The aim of a peak detection circuit is to determine the zero rate of change condition of the input waveform with a minimum amount of jitter. One important property of such a system is that amplitude variations in the input signal can occur without significantly affecting the detection time. However, if noise is present then deviations from the peak occur even for small amplitudes of noise. Furthermore if the detection circuit has a high frequency response then detection can occur even at positions significantly away from the peak *(Fig. 5)*.

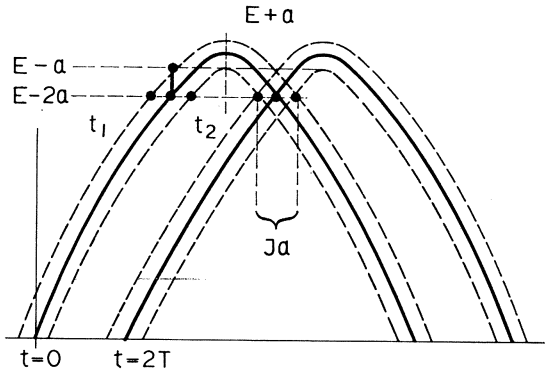

Fig. 10
Description of angle detection

In the angle detection scheme a delayed version of the input signal is subtracted from itself. This can be done quite effectively using a delay line, and detection occurs when the output signal crosses the base line. This in fact will be at a time delayed from the peak. Now any high frequency whose period $T \ll T_d$ (the signal delay) can be biased off by the amplitude of the signal and is not magnified as in the peak detection scheme by its high rate.

To compare the peak detection and angle detection methods consider a system incorporating a peak detection circuit which has a jitter of $\pm T$ about the peak. Now consider that this is replaced by an angle detection circuit with a delay of $2T$. The waveforms are shown in *Fig. 10*. One half cycle of the read signal of amplitude E is shown and the signal delayed by $2T$. Each of these sinusoids has superimposed noise represented as a shift in the sinusoid of $\pm a$ shown by the dotted lines. The movement of the point of zero difference is shown as distance J_a.

Assume

$$E \cos \omega T = E - 2a \tag{1}$$

$$\omega T = \cos^{-1}\left(1 - \frac{2a}{E}\right).$$

Also

$$E \sin \omega t_1 + a = E - 2a$$

$$\omega t_2 = \sin^{-1}\left(1 - \frac{3a}{E}\right)$$

and

$$E \sin \omega t_2 - a = E - 2a$$

$$\omega t_2 = \sin^{-1}\left(1 - \frac{a}{E}\right)$$

but

$$J_a = t_2 - t_1$$

$$J_a = \frac{1}{\omega}\left[\sin^{-1}\left(1 - \frac{a}{E}\right) - \sin^{-1}\left(1 - \frac{3a}{E}\right)\right].$$

Jitter due to Peak detection circuit $= J_p$

where $J_p = 2T$, $J_p = \frac{2}{\omega}\cos^{-1}\left(1 - \frac{2a}{E}\right)$ substituting from (1)

Improvement in reduction in jitter $= J_a/J_p$

$$\frac{J_a}{J_p} = \frac{\sin^{-1}\left(1 - \frac{a}{E}\right) - \sin^{-1}\left(1 - \frac{3a}{E}\right)}{2\cos^{-1}\left(1 - \frac{2a}{E}\right)}.$$

Assume $a = \frac{E}{10}$

$$\frac{J_a}{J_p} = \frac{1}{3.7}.$$

An improvement of approximately four to one which has been confirmed in practice on the magnetic drum system.

In any big system where peak detectors are remote from the actual storage mechanisms so that they can be shared between several switched systems then some noise is unavoidable. In these cases it is important to use angle detection, but in any case there is always this factor of improvement to be gained since an angle detection circuit is no more expensive than peak detection. In practice the longer the delay the bigger the improvement, but the circuit then becomes sensitive to asymmetry which can exist in the waveforms due to the digit patterns stored. The maximum delay is then a compromise between these two conditions.

REFERENCES

[1] KILBURN, T., EDWARDS, D. B. G. LANIGAN, M. J. and SUMNER, F. H., One-Level Storage System, *Trans. I. R. E.*, EC—11, April 1962, p. 223.

[2] KILBURN, T., HOWARTH, D. J., PAYNE, R. B. and SUMNER, F. H., The Manchester University ATLAS Operating System, Part I: Internal Organisation, *The Computer Journal, 4*, Oct. 1961.

[3] ZIMAN, G. C., Maximum Pulse-Packing Densities on Magnetic Tape, *Electronic Engineering, 34*, Aug. 1961, p. 521.

[4] EDWARDS, D. B. G., LANIGAN, M. J. and KILBURN, T., Ferrite Core Memory Systems with Rapid Cycle Times, *Proc. I. E. E., 107* B, Nov. 1960, p. 585.

[5] KILBURN, T., EDWARDS, D. B. G. and ASPINALL, D., A Parallel Arithmetic Unit Using a Saturated Transistor Fast-carry Circuit, *Proc. I. E. E., 107* B, Nov. 1960, p. 573.

The Structure of the Ferranti Atlas Computer

Ferranti Ltd

Computer Department

Enquiries to:

London Computer **Centre**

68 Newman Street, London W1

Telephone MUSeum 5040

List CP. 1
Printed in England
September 1961

No. 1 in a series of Computer System Profiles

Ferranti

1

ATLAS Computer

Ferranti ATLAS Computer

- ★ **Atlas** achieves a break-through in computing technique, bringing speeds of the order of 1,000,000 *completed* instructions per second.
- ★ **Atlas** provides the lowest costs per operation by working at the highest speed.
- ★ **Atlas** puts virtually unlimited storage at the users' disposal.
- ★ **Atlas** automatically allocates storage between different levels in the most efficient way.
- ★ **Atlas** looks like a single level store machine to the user.
- ★ **Atlas** provides comprehensive time-sharing with complete program protection.
- ★ **Atlas** uses a super-speed fixed store for functional sub-routines and peripheral handling routines. This gives exceptional flexibility.
- ★ **Atlas** can handle any configuration of peripheral equipments, with multiple operating consoles.
- ★ **Atlas** brings cheap large scale computing to a new world of computer users.

Storage System Internal Stores

Fixed Store
Access 0·3 micro-seconds, 8192 words (extendable to 262,144 words)

This store is constructed by inserting ferrite rods into a woven wire mesh. It is used for storing functional sub-routines known as 'Extracodes' and fixed programs.

B–Store
Access 0·35 micro-seconds
128 half-words

This store holds indices (modifiers) and has its own accumulator which can operate concurrently with the Main Accumulator.

Main Core Store
Access 2 micro-seconds through 4 or more independent access systems.
16,384 words upwards

The independent access systems permit overlapping of instructions, successive commands being routed through separate systems. In this way overall speeds of the order of 1,000,000 instructions per second are achieved.

Subsidiary Store
Access 2 microseconds

Working space for fixed store routines, locked out from the main store.

V–Store

Data signals and control signals for peripherals.

Drum Store
4 or more 25,576 word drums, access 12 milliseconds maximum

Drum transfers are autonomous, and cause the machine to obey other programmes until the transfer is completed.

The Single Level Main Store

From the users point of view the main store comprising cores and drums together may be regarded as continuous. The Main Store may be considered as made up of blocks of 512 words (pages of core store and sectors of drum store). A relative addressing system is used whereby each block of 512 words carries a unique block number with it. When such a block is held as a page of the core store, the block number is contained in a special register (the page address register) associated with that page. Whenever a reference is made to the main store a comparison is made between the block number of the reference address and the contents of all the page address registers. If the word referred to is in the core store one of the comparisons will succeed and the word will be selected from the corresponding page. If on the other hand no comparison succeeds, a fixed store programme is automatically entered and interchange of blocks between the core store and the drum store is effected, so that the required word becomes available in the core store.
The drum transfer routine in the fixed store optimises the timing of the transfers, using a 'learning' program

RE OF ATLAS

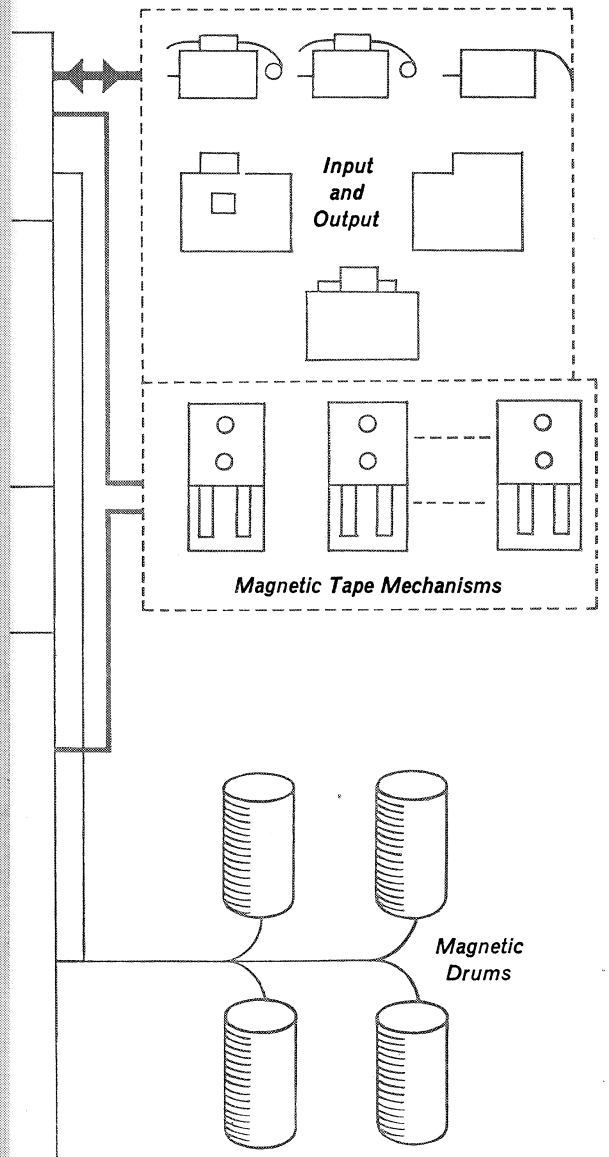

and updates the page address registers and its own records of the sectors used by blocks.

Words and Instructions

Words are 48 bits long.
Instructions occupy a 48-bit word.
Floating point numbers of the form $x8^y$ have an 8-bit signed exponent, and a 40-bit signed mantissa. The octal exponent speeds shifting.
Words may be used to hold 8 6-bit characters, numbered 0 to 7.
Addresses comprise a 20-bit word address, together with 3 bits for a character within a word, or for a half word and character within a half word, and a single bit to distinguish Main Store addresses from other internal addresses.
Two index register numbers B_a and B_m are referred to in an instruction. This permits double modification of arithmetic instructions and two-address indexing instructions.
The instruction format is as follows:

Function	B_a	B_m	Address	Character
10-bits	7-bits	7-bits	21-bits	3-bits

A typical arithmetic operation is:
0320, 51, 52, 1234
'add the floating point number in register 1234 + i + j to the accumulator and round off, where i and j are word addresses in index registers 51, 52'.

Speeds

The times taken by instructions depends very much on their context because of the overlapping of instructions resulting from the multiple access to storage and the independent access to the index registers. Thus as many as three indexing operations may be completed while an arithmetic operation is in progress. As an approximate guide, a floating point addition takes 1·3 microseconds, a multiplication about 3·5 microseconds, the product of two n vectors may be formed in about 10n microseconds and a power series may be summed to n terms in about 7·5n microseconds.

Sequence of Operation

Normally the machine is obeying instructions taken from the main store. The address of the instruction being obeyed is held in the **Main Control** in the special purpose index register B127 in the B-store.
If, however, a complicated instruction requiring, for example, the formation of the logarithm of the number in the accumulator is required, the function digits corresponding to the function 'logarithm' are copied into the **Extracode Control Register** (B126 in the B-store), and the logarithm is computed by the extracode routine in the fixed store, starting at an address in the fixed store corresponding to the address in B126. When the extracode routine is completed, control reverts to the Main Control in B127.

The extracode facility allows the basic instruction code of the machine to be augmented to include about 250 additional functions for elementary functions, input and output conversion and mixed radix conversion. In short all the facilities normally thought of as part of a sub-routine library are available in **Atlas** as extracode functions.

If a peripheral transfer terminates or if any peripheral device requires access to the computer while either main or extracode instructions are being obeyed, control is transferred to a third control register stored in index register B125, known as **Interrupt Control.** All peripheral transfers are initiated by extracode functions. Interrupt control is called in automatically whenever an information transfer (usually of one character, column or line) is required to enable the device to continue at full speed. The transfer is organised by a fixed store program, which passes control back to the interrupted program when the unit of information concerned has been transferred. In the case of the drum and magnetic tape transfers the initiation of the transfer is handled by an interrupt routine, but thereafter the transfer and a program proceed concurrently, the transfer

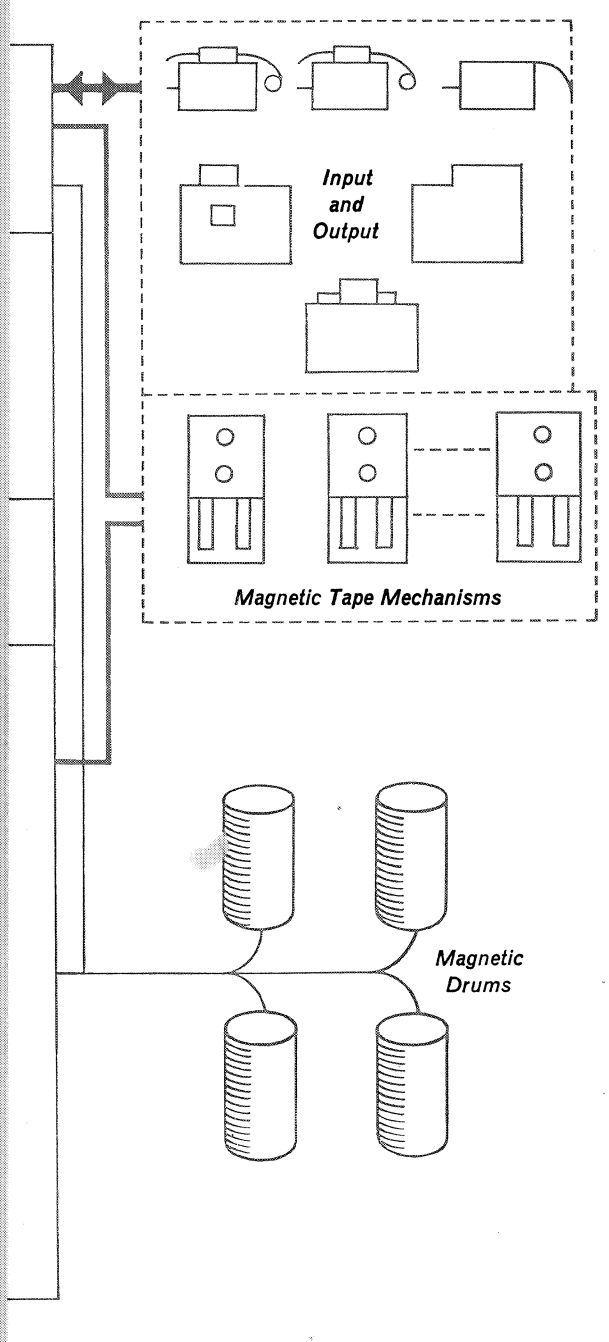

causing the program to hesitate when access to a word in the core store is required by the transfer.

Time-Sharing

The hardware provided for these operations enables **Atlas** to handle any number of programs concurrently, without any possibility that an error in one of them will interfere with any other. The supervisory routine in the fixed store will review the priorities accorded to programs from time to time, and will adjust these in the light of operators instructions, and in the light of experience accumulated by the supervisory routines, so as to ensure the most efficient working of the system as a whole. The standard provision is for four consoles, thus giving four users access to the machine. Each operator may, however, be running more than one program at a time.

The Peripheral System for Atlas

The standard provision for hardware for peripheral equipment for **Atlas** is as shown below. (Further equipment requires additional hardware.)

	Standard Provision
Input Equipment	
Clock	1
ICT Card Reader (600 cards/minute)	4
TR5 Paper Tape Reader (300 characters/second)	12
TR7 Paper Tape Reader (1000 characters/second)	4
Output Equipment	
Teletype Punch (110 characters/second)	12
ICT Card Punch (100 cards/minute)	2
ICT Printer (600 lines/minute)	2
Creed 3000 Punch (300 characters/second)	4
Xeronic Printer	2
Operators Input/Output	
Flexowriters and Teleprinters	16
Magnetic Tape	
Ampex TM2 Units (90,000 characters/second)	32

THE STRUCTURE OF ATLAS