Three-wire mass core store

by **D. J. Morris**, B.Sc., Dipl.Eng., M.Sc., *formerly English Electric Computers Ltd*

The article describes a 16Mbit mass core store with a low cost/bit and a relatively slow cycle time of 8μs. It uses a 3-wire system designed to reduce the cost of the electronic circuits to a minimum.

La matrice à mémoire à noyau de masse de 16M chiffres binaires dont traite cet article est à chiffres binaires d'un coût de revient modique et à temps cyclique relativement lent de 8μs. Elle comporte un système à 3 fils conçu pour réduire au minimum le coût des circuits électroniques.

Der Aufsatz schildert einen 16 bit Grosskernspeicher mit niedrigen Kosten je Bit und einem verhältnismässig langsamen Rechenzyklus von 8μs. Durch das verwendete Dreileitersystem lässt sich der Aufwand für die elektronischen Schaltungen auf ein Minimum herabsetzen.

A MASS CORE STORE has been defined as a core store with a capacity of more than 5 Mbits. It must be able to store information more cheaply than the main store and must also have a faster cycle time and transfer-rate than any existing backing store. However, the function of the mass store in a computer is not yet well defined, and there are two distinct approaches to its use in a computer system. The first regards it as an extended main store, having the speed of operation as its main design criterion, with the system cost only of secondary importance. The second approach regards the mass store as a peripheral backing store, where the main criterion is the cost, with the speed only a secondary requirement. These two different approaches will also affect the design of the interface between the central processor and the mass store. In fact, the mass store may be used either as an extended main store or as a backing store, depending on the logical connexions to the computer.

The mass core store described in this article was designed as a peripheral store with the aim of achieving a low cost/bit, and a cycle time of less than 10μs. Emphasis was put on low cost rather than high speed.

The capacity of the store described here is 16 772 216 bits, i.e. 524,288 words each of 32 bits. The initial tests carried out on a small model of this store indicate that it can operate successfully at a cycle time of 8μs but it is believed that this can be decreased to about 5μs.

Design philosophy

Although core planes are used in both main stores and mass stores, their design philosophy is different. In mass stores, owing to the large number of bits being manipulated, emphasis is on the cost of the actual core stacks, while in main stores the electronic circuits take the major portion of the overall cost. It has been established that in a main store the cost of the core stacks is about 40-50% of the cost of the whole system, and that of the electronic circuits is about 50-60%. In the mass store this portion is reversed, and the core stacks represent 60-70% of the total cost of the system, and the electronic circuit only 30-40%.

The cheapest core store plane is one that has only vertical and horizontal wires, i.e. only X and Y lines. However, a store using a 2-wire system suffers from various technical drawbacks which make the design difficult to translate into practice. Furthermore, it makes the electronic circuits complicated and expensive. By using a 3-wire system, but still keeping the principle of only vertical and horizontal wires, these difficulties can be reduced and an overall cheaper system may be obtained. The low cost of the 3-wire system is based on the fact that the cost of the third wire increases the cost of the stack of cores by only 5%, but reduces the cost of the electronic circuits by over 5%, relative to the cost of the 2-wire system. However, this reduction in cost is obtained at the expense of the cycle time of the system.

The principle of the three-wire system is to drive as many cores as possible in series and thus to have common electronic circuits for a large part of the store. This requirement is limited by the availability of fast power transistors which can supply the required current at reasonably fast risetimes and can also contend with the large back e.m.f. presented by the core load. The latter is restricted by three factors: the resistance, capacitance and inductance of the wire. In the system described here,

X = number of lines across the width of the plane, i.e. in the X-direction
Y = number of lines along the length of the plane, i.e. in the Y-direction
$X \cdot Y$ = plane size
N = number of planes in the stack
n = number of bits per word
n = basic number of bits per word
K = maximum number of cores that can be driven in series
$N \cdot X \cdot Y/n$ = number of words in the stack
$N \cdot X/n$ = number of word-groups in the stack
K/Y = Maximum basic number of bits per word

the driving of a line threading 65 000 cores with an ohmic wire resistance of 185Ω has been achieved.

Basic organization

The 3-wire store is organized in a multi-planar stack, as in Fig. 1, which differs from the conventional 3-dimensional system by having the full word stored on a given plane; whereas in a 3-dimensional organization, each bit of a given word is stored on a separate plane.

In this 3-wire system, only three selection lines are used, the Y, the word-group and the sense-inhibit. A word-group is formed by connecting in series n consecutive X-lines of the same plane, where n is the number of bits per word. Every Y-line and sense-inhibit line is wired in series through all the planes, whereas separate word-group lines are use for each plane. The word-group lines are parallel to the sense-inhibit lines but are perpendicular to the Y-lines.

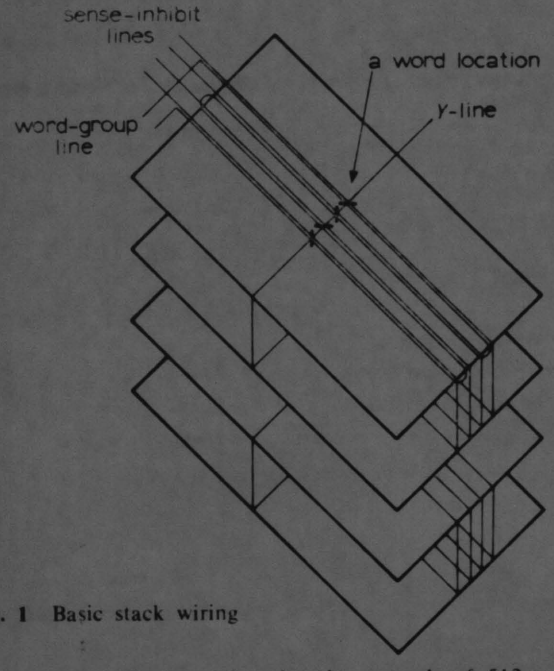

Fig. 1 Basic stack wiring

The system designed has 16 planes each of 512 × 2048 cores which can be adjusted as required. The word-group line is wired in series through 32 lines (each of 2048 cores), giving a 32-bit word length, but this too can be modified according to circumstances. In Fig. 1, for demonstration purposes, the word consists of four cores, i.e. only four X-lines are driven in series for each word-group.

Three-wire organization

As mentioned before, the 3-wire system uses coincident currents for both the READ and WRITE operations. The stack consists of N planes each of XY cores, where X represents the number of lines across the width of the plane and Y represents the number of lines along the length of the plane. There are X/n word-group lines on each plane and NX/n word-group lines in the stack, where n is the number of bits per word. The maximum basic number of bits per word will depend on the maximum number of cores that can be driven in series; hence it will be $\bar{n} = K/Y$, where K is number of cores that can be driven in series. If the required number of bits per word n, is larger than the basic number of bits per word $\bar{n}$, then two word-groups will have to be driven in parallel and operated together. The number of cores driven off each Y-line is NX and the number of cores driven by the inhibit lines is NY. As will be explained later, two sense-inhibit lines are operated in parallel, making the number of cores on that line $2NY$.

Reading is accomplished in two sequential selection stages and the output of the selected word appears simultaneously on the corresponding sense-inhibit lines. The word is stored at the intersection of the word-group and Y-lines, having the number of bits read (i.e. word size) depending on the number of times these lines intersect.

In the reading operation, a long-duration current pulse (of half the magnitude required to switch a core) is first driven through the word-group line. This will cause all the cores in this line to move magnetically to just below the knee of the hysteresis loop; i.e. as this current is not of sufficient magnitude to switch the cores it will only cause 'elastic switching', or what is commonly termed 'disturb'. The output from the disturbed cores is usually small compared with the actual remanent switching of cores. However, since the sense lines are parallel to the word-group lines, the cumulative voltage pulse induced by all the disturbed cores into the sense lines, is very large compared with the remanent output pulse from a single core. Thus the cores storing the required information cannot be selected for reading until the disturbed output pulse has fully decayed, i.e. the elastic switching has been completed. Only then can a second current pulse of shorter duration and of the same amplitude be applied to the selected Y-line. As this stage both current pulses are applied, thus only the cores at the intersection of the two drive lines will receive the full required current and be able to switch. Although all the other cores on the Y-line will get the half-magnitude pulses, the 'disturbed' output pulses from the majority of these cores will not be induced in the selected sense lines.

The word-group line drives nY cores in series causing elastic switching in all of them. Each sense line, which is parallel to the driven word-group line, will be induced from the cumulative output from Y cores. When the Y-pulse is applied, it will tend to switch n cores which are situated at the intersection with the selected word-group line. At this instance, each sense line will be induced by the remanent switching from only one core.

This operation may be more easily explained with the aid of the model shown in Fig. 2, which represents two planes each with 16 cores. This model can store 16 words each of two bits: 1 and 2, 3 and 4, 5 and 6 etc. If the word stored in cores 1 and 2 is to be read out, then word-group 1 and Y_1 lines are selected. The long duration pulse is first applied to word-group 1, which will disturb cores 1, 2, 5, 6, 9, 10, 13 and 14, causing elastic switching, while all the other cores in this model will remain unaffected. The disturbed output from cores 1, 5, 9 and 13 will be induced in sense line D_{11} and that of cores 2, 6, 10 and 14 in sense line D_{12}. Once these disturbed voltage pulses, due to the elastic switching, have decayed, Y_1 is selected, which will switch cores 1 and 2 fully (if a 1 had been stored) and the output from these will be induced in the respective sense lines D_{11} and D_{12}. This Y-drive pulse will also disturb cores 3, 4, 20, 19, 18 and 17, but the output from the first four will not be induced into the selected sense line.

The outputs from cores 17 and 18 in the model will, however, be induced in the respective sense lines at the same time and in the same direction as the signals from cores 1 and 2. Means for eliminating these unwanted output pulses will be explained later.

In a full multiplane stack, each sense line will be first induced with the disturbed output from Y cores, where Y is the number of core along the length of a plane. Then when the second pulse is applied, it will be induced with the signal output from one fully switched core, and $N-1$ partially switched cores (disturbed output), where N is the number of the planes in the stack. This is due to the fact that each Y and sense line intersect once on each place. The problem of eliminating these unwanted

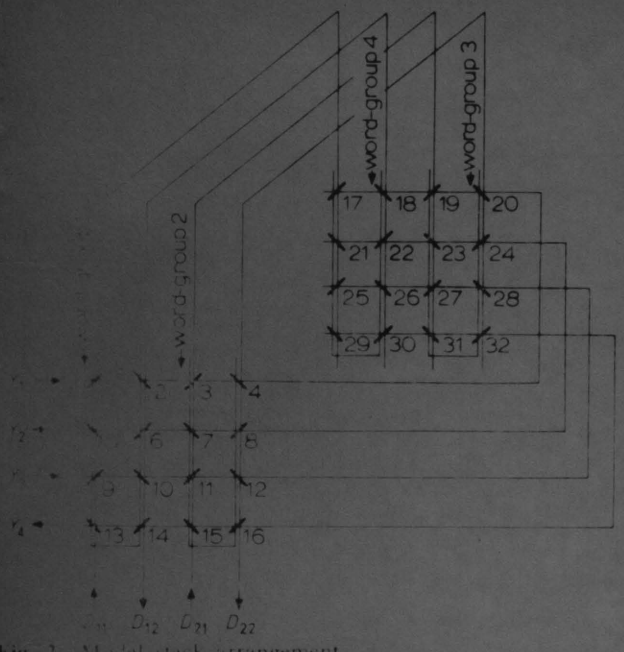

Fig. 2 Model stack arrangement

disturbed outputs is easier than in a conventional 3-dimensional organization where there are $X + Y - 2$ disturbed cores on each sense line for every READ operation.

The selection of a word for the writing operation is accomplished in a similar manner to the read operation, i.e. by the selection of a word-group line and a Y-line. The writing is carried out by the coincidence of two current pulses, both of short duration, applied to these two selected lines, but with opposite polarities to those of the READ current pulses. In this way a 1 will be inserted in all the cores at the intersections of the two selected drive lines. In order to write a 0 into a core, the specific core must be inhibited during the write operation, to prevent it from switching. This is done by applying an inhibit pulse to the sense line of the appropriate digit in the opposite direction to the word-group WRITE current in the selected core. The arrows on the various drive lines of the model shown in Fig. 2 indicate the direction of the three current pulses during the WRITE operation.

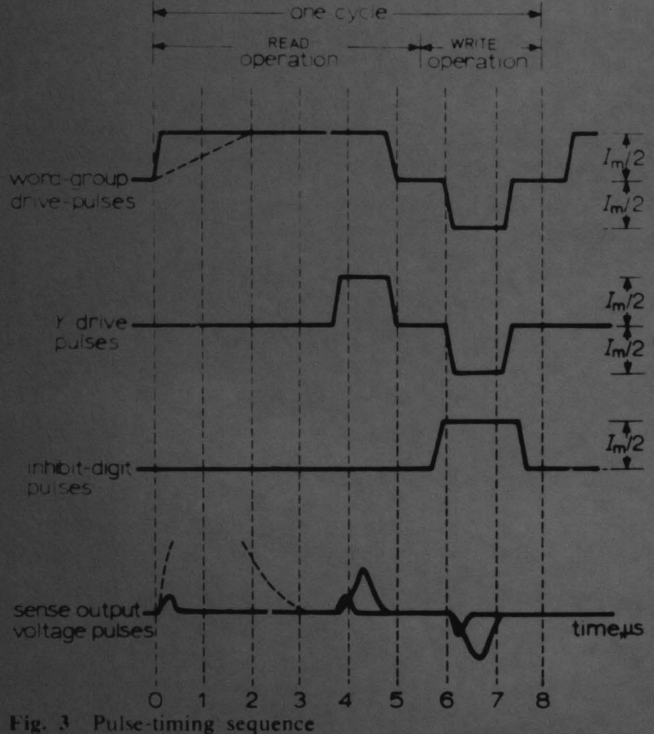

Fig. 3 Pulse-timing sequence

The digit lines are used for both inhibiting the cores during the WRITE operation and sensing when reading the information stored in the cores. Although these two operations apply to the same wire, they do not affect each other because they occur at different times in the cycle.

The timing sequence of the various pulses applied to the store, for both the READ and WRITE operations is shown in Fig. 3.

The Y-driver and selection matrix

The Y-selection lines are driven through all the planes and hence each line drives NX cores. The number of the Y-lines depends on the number of lines along the length of the planes. As the store is designed to have $Y > X$, the number of the Y-lines is rather large, and hence the selection matrix with its associated switches could be rather expensive. By using double-level selection, the cost of the selection unit is greatly reduced.

If a conventional diode selection matrix is used, a matrix size of $4Y$ diodes is required, using four steering diodes per drive line. This matrix will require $4\sqrt{Y}$ switching transistors. The switching circuits are designed to pass current pulses in either direction in the selected line, hence four transistors are required for each line two for the READ operation and two for the WRITE operation.

Owing to the large core load on each line, the back e.m.f. presented by each line is rather high, and thus expensive high-speed high-power transistors are required. To reduce the number of these switching transistors for the selection of the Y-lines, a 2-level selection is used, having two transformer matrixes selecting the diode matrix. This theoretically reduces the number of switching transistors to $8\sqrt{Y}$, although in practice (as explained later) some of these switches can be common for reading and writing, thus further reducing the number of transistors to $6\sqrt{Y}$. This double-level system, however, calls for $4\sqrt{Y}$ transformers which are considered to be cheaper than the switching transistors.

The double-level selection circuit for the selection of the Y-lines for both reading and writing is shown in Fig. 4. Four transformer matrixes are required two for the READ operation and two for the WRITE operation. Each transformer in a matrix is selected by two transistors, one working as a current switch to select the columns and the other as a voltage-level switch to select the rows. The voltage-level switches may be used for both READ and WRITE transformers, hence they are common to two transformer matrixes.

For a READ or WRITE operation, each Y-line is selected by a combination of four transistors which select two transformers. The steering diodes are reverse-biased by the positive and negative voltage levels of the transformer secondaries. Two transformers are selected for each operation and the voltage pulse induced in their secondaries will forward-bias the diodes, causing the required current to flow in the selected line. For example, line Y_6 in Fig. 4 is selected for the WRITE operation by transformer 3, which was selected by VT_1 and VT_3, and by transformer 16, which was selected by VT_8 and VT_{12}. When selecting the same line (Y_6), for the READ operation, transformer 7 (selected by VT_5 and VT_1) and transformer 12 (selected by VT_8 and VT_{10}) are seleceted. Transistors VT_1 and VT_2, which are voltage-level switches, are selected for both READ and WRITE operations.

As described before, this system consists of 2048 Y-lines which will require a diode matrix of 64×32. If the conventional selection method is used, $2(64 \times 32) = 192$ switching transistors are required. By employing the double-level selection, four transformer matrixes are used, two of 8×8 and two of 8×4. This requires only 40 transistors, i.e. $8 + 8 = 16$ for the voltage level switches and $2(4 + 8) = 24$ for the current switches. A

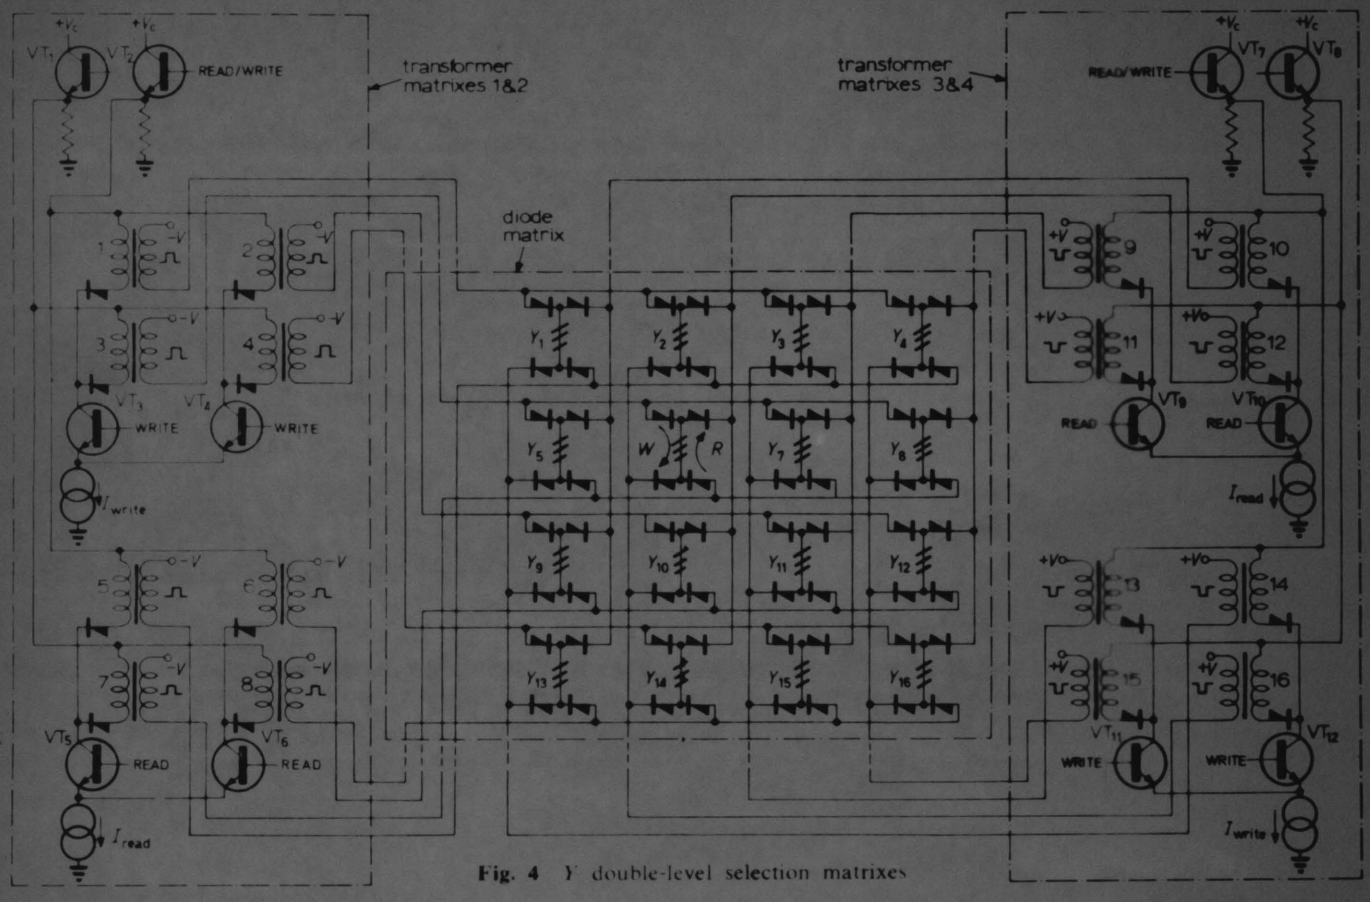

Fig. 4 Y double-level selection matrixes

reduction from 192 to 40 transistors is achieved by the double-level selection; although 192 transformers are also required in these matrixes, they are cheaper than the fast high-power transistors which would otherwise have to be used.

The Y current pulses in the store, with a load of $8K$ cores in the line, are shown in Fig. 5. This load represents 16 planes with a plane width of 512 cores. A resistance was added in parallel to the Y-line to prevent overshoot and ringing on the current waveforms; this resistance is not shown in Fig. 4.

Word-group driver and selection matrix

There are NV/n word-group lines in the stack and each line drives nY cores. The transistor circuits driving these cores must supply current pulses in the lines, a long-duration READ pulse in one direction and a short-duration WRITE pulse in the other direction. As the READ-current pulses are driven for about 60-70% of the cycle time, the transistors supplying these currents must be capable of dissipating a considerable amount of heat.

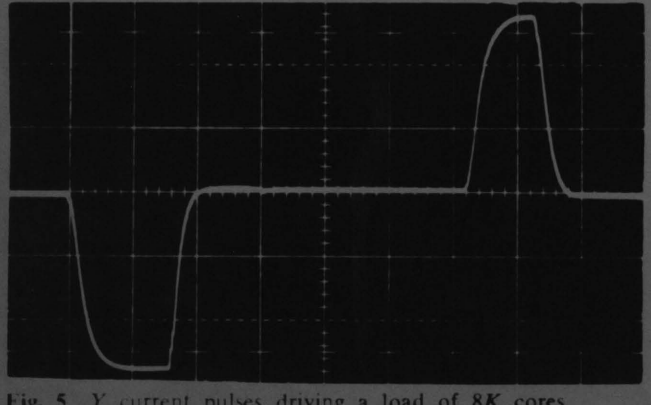

Fig. 5 Y current pulses driving a load of $8K$ cores (100mA/div., 1μs/div.)

In this system, there are only 256 word-group lines and each line has a load of 65 536 cores. As only a relatively small number of lines are involved and the switching transistors are required to handle a considerable amount of power, it is not practical to use a transformer selection matrix, as with the Y-lines; the conventional selection matrix of 16 × 16 diodes is used.

It has been found that the ripple and overshoot of the READ word-group current pulse can be improved by limiting the current with added resistance rather than by clamping the current at the required value. However, the

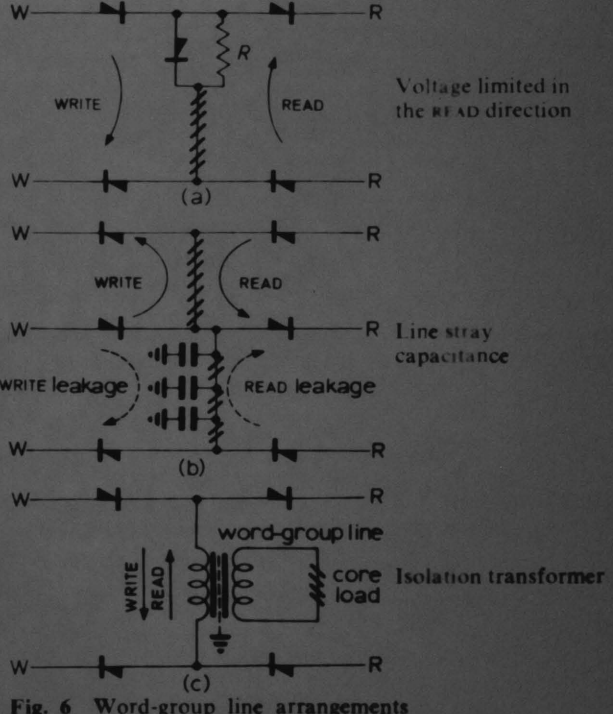

Fig. 6 Word-group line arrangements

current in the WRITE direction is still limited only by the current driver, so as not to increase the WRITE risetime, which is done by bypassing the limiting resistance with a diode, Fig. 6(a).

The current waveform in a word-group line with a full load of 65K cores is shown in Fig. 7. Detail (a) shows the current in a line with no limiting resistors, (b) shows the current in the same line, but with a limiting resistance of 66Ω in the READ direction, and (c) shows the current in a selected core, reading and writing a 1, with full core load on both the Y and word-group lines.

Owing to the large number of cores driven in series and the long length of wire, the capacitance presented by the line (to virtual earth) is rather high. This does not cause trouble when driving a single line, but creates difficulties when using a selection matrix. Owing to this capacitance, a small amount of drive current is fed into an unselected line, as can be seen from Fig. 6(b), where the top line is selected while the bottom one is unselected. Although the two lower diodes are biased off, a leakage current flows into the unselected line during the risetime

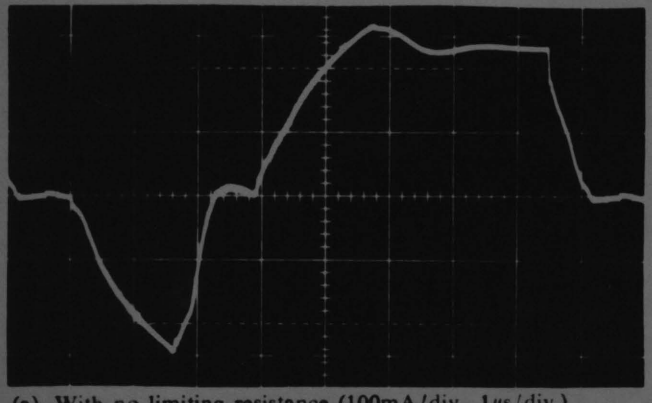

(a) With no limiting resistance (100mA/div., 1μs/div.)

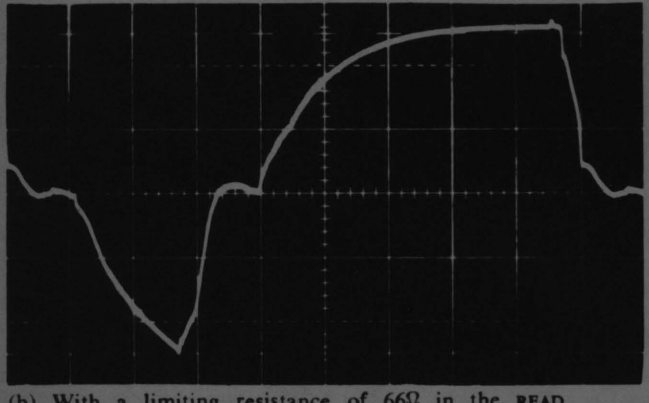

(b) With a limiting resistance of 66Ω in the READ direction (100mA/div., 1μs/div.)

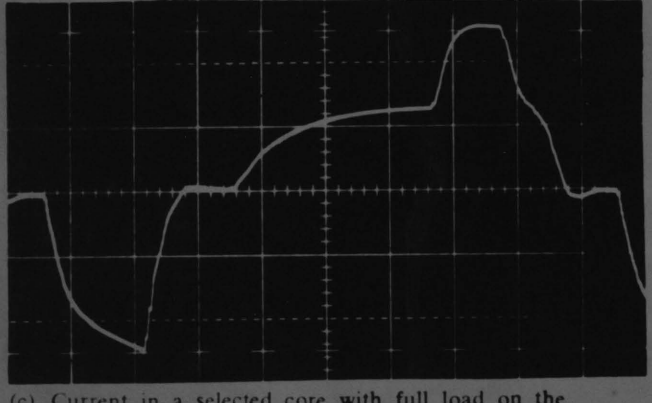

(c) Current in a selected core with full load on the Y and word-group lines (200mA/div., 1μ/div.)

Fig. 7 Current waveforms in the word-group line with a full load of 65K cores

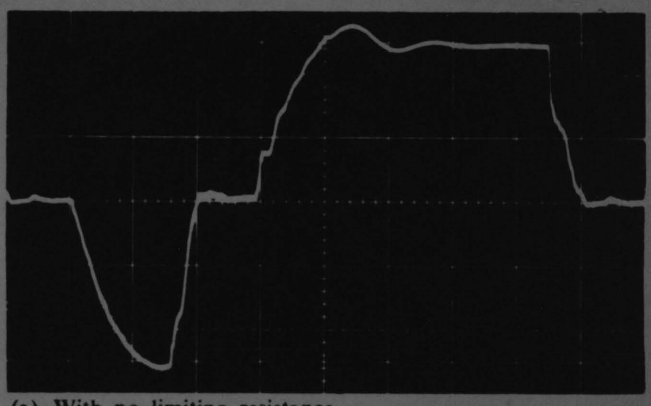

(a) With no limiting resistance

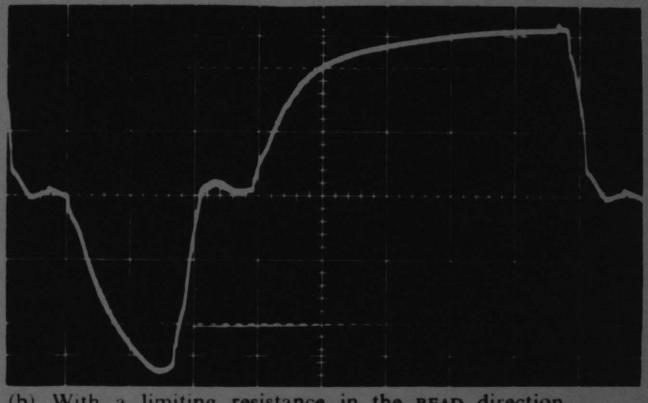

(b) With a limiting resistance in the READ direction

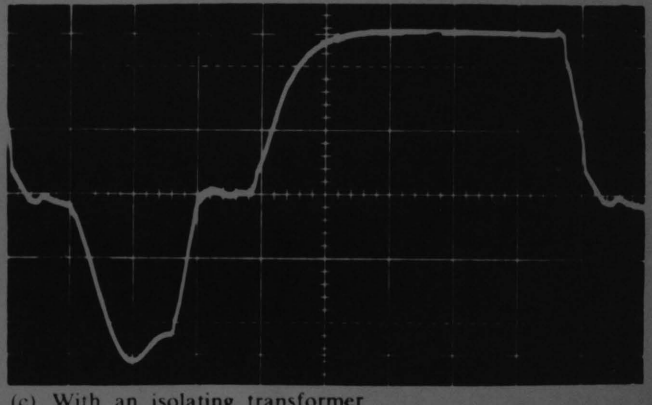

(c) With an isolating transformer

Fig. 8 Current waveforms in a word-group line with a load of 32K cores (100mA/div., 1μ/div.)

of the pulse, charging the stray capacitance to earth. In order to reduce the effect of this capacitance and prevent the leakage current flowing in the unselected line, the core lines are isolated. This may be done by coupling the lines to the diode matrix with a transformer having an earthed shield between the primary and secondary windings, as in Fig. 6(c). This arrangement is not necessary in the Y-lines, as the drivers are already isolated from the cores by the transformer matrix.

The testing of the isolating transformer is incomplete but results are available from testing a word-group line with a core load of 32K cores, shown in Fig. 8. Detail (a) shows the current with no limiting resistance and (b) with a limiting resistance of 154Ω in the read direction, and (c) shows the same current as (b) but with an isolated transformer. Comparing (a) and (b), it can be seen that the ripple and overshoot of the READ pulse is reduced. Comparing (b) and (c), that the isolated transformer also improves the rise time of the current pulses, suggesting that, by using an isolated transformer, faster cycle times could be achieved.

(To be continued)