6889
Memorandum M-1495

Digital Computer Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts

SUBJECT: WHIRLWIND II MEETING OF APRIL 25, 1952

To: Whirlwind II Planning Group

From: N.H. Taylor, R.P. Mayer, and W. Papian

Date: May 21, 1952

Members Present:

I. Aronson	W. Hosier	W. Ogden
W. Attridge	J. Jacobs	K. Olsen
P. Baltzer	R. Jeffrey	R. Pacl
G. Briggs	W. Linvill	W. Papian
D. Brown	R. Mayer	C. Schultz
D. Eckl	J. McCusker	N. Taylor
R. Gerhardt	J. Mitchell	J. Woolf
A. Guditz	R. Nelson	

N. Taylor pointed out that in attempting to design a new computer such as WWII, there are a number of engineering problems which arise and which should be called to everyone's attention. The magnetics and transistor groups are running into engineering problems which will affect the logical design of the computer. The logical design group must become familiar with these problems. In order to help in this direction, W. Papian presented some of the problems which are arising in connection with the magnetic memory and the magnetic switch.

W. Papian first listed some of the general things that are being worked on. These included the metallic array (which contains 256 Moly-Permalloy cores), the ceramic array and switch Model II, and the ceramic array and switch Model I which will be described and discussed today.

The Model I array has 256 cores arranged in a 16 x 16 core array. The array itself is 4" by 4", and around this square the switching cores are arranged. The cores of both the memory and the

switch are shaped like doughnuts. The cores of the memory itself are rather small, are called "Cheerios", and are made of Ferramic 1118. The "Cheerios" are single-turn driven; the single turn consists of a straight wire on which the cores are strung. The cores of the switch are somewhat larger and each one has six twenty-turn windings and one two-turn secondary winding.

The present performance, which is outlined below, is based on the 8 x 8 corner of the array now in operation; the full 16 x 16 array has not yet been operated.

Very generally speaking, class C pentodes (6CD6's) bias and drive the switch cores with a 300 ma., one-microsecond, square pulse which produces on the selected secondary a $1\frac{1}{4}$-ampere square pulse to the memory cores. The selected memory core delivers to the sensing line a positive or a negative $\frac{1}{3}$-volt pulse if it contained a "1" or a 0.1-volt pulse if it contained a "0". This signal is fed through a pair of 5-to-1 step-up transformers and diodes connected in a full-wave rectifier circuit. The voltage output of the full-wave rectifier is 2 volts (single polarity), which is then fed to an amplifier composed of two 7AD7 tubes. The resulting 20-volt signal is fed to a Burroughs gate panel. (It was suggested that perhaps one transistor could be used instead of the two 7AD7's).

To be a little more specific about the performance, consider the sketch shown in Figure 1. This sketch shows the output waveform that is likely to appear on the sensing line as the result of selecting a given core. The waveform, of course, is not as thick as that which is shown. Rather, the thickness indicates the range of values within which a single core is likely to appear. This is the waveform that can be seen on an oscilloscope when the cores of the array are sensed in quick succession, with the addition that an approximate mirror image of the waveform appears below the axis due to the fact that the output pulse may be positive or negative.

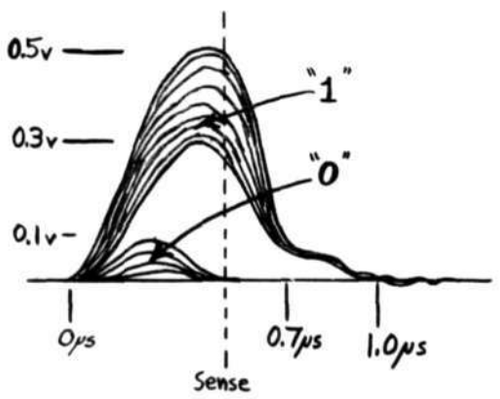

Figure 1.

Notice that there is a factor of about 3 to 1 between the peak amplitude of the smallest "1" signal and that of the largest "0". This signal-to-noise ratio can be improved by sensing the out-put pulse at the indicated time when the "1" signal has dropped off hardly at all, while the "0" signal has dropped off considerably. A 10-to-1 ratio can be obtained in this manner.

The "0" output signal is due to the fact that the square hysteresis loop is not exactly flat but is somewhat slanted. This means that reading-out a "0" produces a small output pulse, but it also means that a non-selected core (pulsed with $I/_2$) will produce an output pulse of the same order of magnitude. Since the slope of the hysteresis loop is about the same in either condition of saturation, the "non-selected" pulse will have about the same amplitude regardless of whether the core contains a "1" or a "0". If all the non-selected outputs were allowed to accumulate, they could very easily override the signal from the selected core.

In order to avoid accumulation of these signals, they are made to cancel each other in the following manner: The sensing winding is wound in opposite directions through adjacent cores so that the undesirable signal from adjacent cores will be canceled out. This means, of course, that the desired signal from the selected core will be positive in one case and negative in another. The design of the storage array and sensing winding happens to be such that there are always two non-selected cores (one on each axis) whose signals do not cancel. Instead, these two signals subtract from the signal from the selected core, so that the signal on the sensing winding is the result of the selected core and two non-selected cores.

It should be pointed out that the unwanted signals will be exactly canceled only if the amplitudes and shapes from the individual unwanted cores are identical. The degree of cancellation is related to the uniformity of the cores on the array. With the 8 x 8 array now being tested, cancellation is very good. In wondering whether it will be as good in the 16 x 16 array, it is interesting to notice that the statistical percentage will be better, but the absolute deviation could be greater. It is possible to use separate sensing windings for different sections of the array. Another alternative is to pre-test the cores and hand-pick those which are sufficiently similar.

The temperature characteristics of the cores may also cause deviations which could prevent effective cancellation, but

it is hoped that forced-air cooling(or, at worst, an oil bath) will be sufficient to prevent such undesirable effects. It is believed that the value of the temperature will not become as objectionable in this respect as the difference in temperature from one core to another.

It was mentioned above that a 1-microsecond pulse is applied. Actually, for safe operation, a 1½-microsecond pulse is applied to the tubes which drive the switch cores. A ½-microsecond interval is allowed between one pulse and another in order to let transients die out.

It is necessary to rewrite the information in a core after each read-out, as shown in the sketch of Figure 2A. Thus a read and rewrite requires a total of 4 microseconds before the next one can start. This method implies that a Z axis will be used to inhibit a rewrite if a zero is to be rewritten. An alternative way of preventing the rewrite of a "1" is shown in Fig. 2B. The "x" and "y" switch cores are cleared separately so that no memory core is actually selected. This requires a separate Y-selection switch for each digit column. Using this method, a total of six microseconds is required for the read-write before another one can start. The cores have never been tested with another operation starting immediately after the six microseconds, but investigation of the waveforms indicates that it could be done.

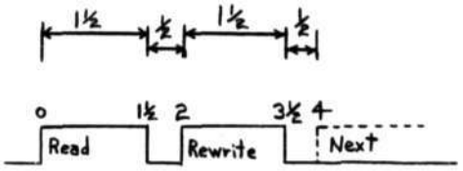

Figure 2A.

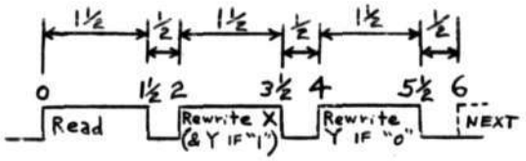

Figure 2B.

Figure 3 shows waveforms for the grid of the selection switch driver, the current supplied to the selection switch cores, the resultant current supplied to the memory array, and the resultant signal on the sensing line. It should be noticed that a total of two microseconds is required for allowing the switch to set up briefly and for the read pulse itself to be completed. This is shown in the sketches of Figure 2 as the $1\frac{1}{2}$ microseconds for the read, preceded by the $\frac{1}{2}$-microsecond wait since the previous operation.

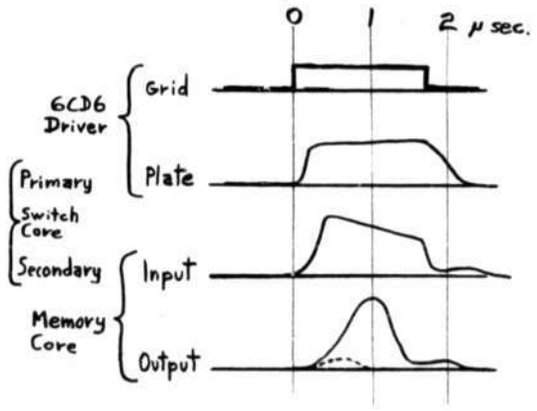

Figure 3

Figure 3 shows that the output from memory can be sensed 1 microsecond after reading the address into the switch drivers. A conservative rough estimate is that it will take 1 to 2 microseconds between the time the address flip-flops are set up and the time the flip-flops on the output of memory are set up.

It is desirable to decrease the amount of time required for the storage operation, but in some cases it does not pay to reduce the rewrite time because this can often be masked by other required operations, such as reading-out of some other memory array. On the other hand, it might be desirable to eliminate the necessity for two separate rewrites, as in Fig. 2B. Thus if the inhibit winding is used, 4 microseconds instead of 6 will be required for the complete read-write. Perhaps it is possible to reduce the $1\frac{1}{2}$-microsecond pulses to 1-microsecond pulses, which would reduce the total time to 3 microseconds. It is interesting to notice that the metallic array, using vacuum tube drivers instead of a magnetic switch, not using the double rewrite and using 3:1 selection, can perform the complete read-write in 8 microseconds. With this arrangement it probably takes 2 microseconds minimum to get the output from the memory after the start of the read pulse.

Discussion by D. Brown, W. Papian, and N. Taylor brought out the suggestion that we should try to design the logic of the computer with plenty of time for components to function, and if this is not adequate, then we should try to speed up the components. In thinking about the speed of the components, we should use figures

which represent the operation of the equipment at the moment, and not figures which look as though they should be obtainable. However, the logical design must be kept flexible enough so that advantage can be taken of any improvements in component speed. Apparently the figures to use - as far as this discussion is concerned - are: 6 microseconds total read-write time and 1½ microseconds access time (from setting up the switch control flip-flops to setting up flip-flops according to the number read out from storage).

W. Papian then reviewed briefly the history of the recent developments and pointed out some of the problems that were encountered and the cures that were applied. It was desirable to develop the memory array as quickly as possible. First the magnetic switch waveforms without loading were observed, and then with cores in the output of a single switch, and finally with cores on two intersecting switch lines. In this arrangement the sensing winding output waveforms were observed for selected and non-selected cores. Each core of the array was checked individually in this fashion. The rather poor results were examined carefully; they indicated that there had been a mistake in the geometry of the sensing winding. After correcting this the same tests were run over about one-third of the 256-core array; results were much better.

Since each switch core bias winding eventually goes to the plate of a triode driver, a voltage will be applied to the plate from the winding when the core is switched. When this happens, the driver tends to act as a clipping diode and tends to prevent the core from switching completely. This was corrected temporarily by placing a resistance in the line from B+ to the winding in order to reduce the plate voltage on the drivers.

A single line from the plate of a driver runs to B+ through a number of windings, each with some capacitance to ground. This arrangement looks as though it should act like a delay line. Such an effect would not be tolerable in the memory array because it would mean that the coincident currents arrive at the selected core at different times, depending on where in the array the core is located. Crude early experiments verified the existence of this delay-line effect but also indicated that the switching action of a given switch core was delayed by about the same amount of time regardless of which end of the delay line it occupied. This anomalous, but fortunate, behavior was documented and temporarily bypassed.

Up to this point, the full array had never been used, but only 4-core by 4-core sections of the array. By testing various 4 by 4 sections over the entire array, an area was found which operated better than other sections. This area was then expanded to include 8 by 8 cores. At first the sensing was on a straight amplitude-ratio basis, but this was later changed to a single-time sense (as shown in Fig. 1), which worked much better.

Output "1"/"0" ratios were still rather poor, and it was felt that this was due to differences in current waveshapes and amplitudes on the different coordinate lines of the array. It was felt that these bad waveshapes were probably due to clipping in the bias drivers, so they were changed to pentodes. This improved the shapes considerably but did not cure the differences in amplitudes. The differences in amplitudes were corrected by changing the number of turns on the driving cores after experimenting to find the relationship between the number of turns and the current output. The waveshapes were still further improved by juggling the components in the driver circuits. The resulting waveshape on the switch core secondary looks as shown in Fig. 3.

Some rough marginal tests were performed on the memory, as follows: The currents driving the selection switch were set to the "middle" of the reliability range, the sensing panel gain and bias levels, and the various timing intervals, were set to "optimum" values. Then the array was made to operate by continuously reading and rewriting all the cores in a static pattern (so that there were 14 disturbances between separate operations of a single core). It was then found that the current driving the selection switch could be varied plus or minus 15% without destroying the stored pattern. It is interesting to notice that the present power supplies are not too stable, so that this percentage looks quite encouraging. If the pattern is made to cycle around, then it appears that about the same margin is obtainable. It is possible to cycle information in the cores all day except for the trouble caused by very large transients in the power supplies.

We now feel that we are ready to work toward a 16 x 16 array. The switch cores to be used in the array will be tested individually, either on the array itself or on a single core tester if a tester seems more convenient. (See minutes of the next meeting.) After the 16 x 16 array is running, we should try the "Z axis inhibit" method of writing, which will shorten the read-write time and give us experience driving as many as 256 cores from one driver.

Rollin P. Mayer

Rollin P. Mayer

William N. Papian

William N. Papian

Norman H. Taylor

Norman H. Taylor

RPM:WNP:NHT/cs