

THE CONSTRUCTION AND OPERATION OF THE MANCHESTER UNIVERSITY COMPUTER

By B. W. POLLARD, M.A., Graduate, and K. LONSDALE, B.Sc.Tech.

(The paper was first received 27th October, 1952, and in revised form 2nd February, 1953. Proofs were made available to the public 4th April, 1953, and the paper was read before the MEASUREMENTS SECTION 14th April, 1953.)

SUMMARY

Details are given of the design and construction of the new universal high-speed digital computing machine now working at the Computing Machine Laboratory, Manchester University.

The constructional techniques adopted for the electronic equipment, the cathode-ray-tube storage units and the magnetic storage drum are described, and an outline is given of the design procedures adopted for the electronic circuits. The techniques used to obtain reliability and ease of maintenance are particularly stressed. The development of suitable maintenance procedures for the operation of the computer continuously over periods of five days in each week is discussed, and details are given of the performance of the computer over a period of some 61 weeks.

(1) INTRODUCTION

Previous papers^{1, 2, 3} have described the research programme on high-speed digital computers carried out at Manchester University under the general direction of Prof. F. C. Williams. In particular,³ the logical design and the facilities of the computer now working at the Computing Machine Laboratory, Manchester University, have been described in detail, and it is the purpose of the paper to describe the construction and operation of the computer, which embodies the results of the experience gained during the course of the research programme.

The design of the new machine was started during the summer of 1949. By December, 1949, all the major details of the physical layout had been decided, and the logical design was completed. The construction was started, and the detailed circuit design was undertaken. Design and construction, and the testing of completed sections of the computer, therefore proceeded together, and by February, 1951, virtually all the machine had been completed and tested. It was then installed at Manchester University, where further testing was undertaken, and the machine was brought into operation at the beginning of July, 1951.

(2) LAYOUT AND CONSTRUCTION OF THE COMPUTER

When deciding upon the techniques to be used in the construction of a large electronic computer, the designer is freed from many of the requirements normally associated with electronic equipments and is able to concentrate entirely upon the two major requirements of reliability and ease of maintenance.

The factors controlling reliability should be the performance of the valves and electronic components, and the object of the mechanical design should be to ensure, so far as possible, that there are no other sources of failure and that the valves and components are operating under ideal physical conditions. The mechanical design must therefore provide a rigid base for the mounting of valve circuits and their associated components, and allow for the reliable electrical interconnection of the large number of circuits which comprise a computer.

For most of the circuits a standard chassis (see Fig. 1) has




Fig. 1.—A standard chassis as used in the computer, showing the positions for mounting the eight B9G type valves, the B3A diode bases, the filament transformer in the centre of the chassis, and the tag rails for mounting the components.

been used which allows a maximum of 8 pentodes, 27 diodes, 66 components and a filament transformer to be mounted, and which has 28 interconnection terminals—14 at each end. All components, other than those which require to be wired directly on the valve base, are mounted on the tag rails, which are of special design. Each tag is held in a moulded Bakelite block; these blocks slide into steel rails, and the blocks are held apart by small moulded projections situated on the block so that they are under the bent-over section of the rail. Thus any leakage path is to earth rather than to the adjacent tag. Furthermore, the breakdown voltage to earth from any tag is 5 kV, whilst the maximum voltage used on these tags is 300 volts with respect to earth. The lead-through terminal insulation consists of $\frac{1}{4}$ -in diameter Tufnol rod, and the terminals are crimped into the end-plates, again ensuring that the leakage path is to earth. The entire chassis is constructed from 0.048-in steel plate, cadmium plated, and is therefore extremely rigid, even before mounting on the hinged frames. The wedge shape of the chassis allows easy access to the valve bases and the inner ends of the components without increasing very much its superficial area, and the fixing holes of the chassis are dimensioned so that it can be mounted on standard Post Office type racks, for testing and experimental purposes. As the chassis are mounted in a vertical plane, there is effectively a metal sheet between the valves and the components. Hence the temperature of the components is not affected by the much larger dissipation of heat from the valves. The open type of construction ensures that no hot spots can occur and result in an unsuspected over-rating of components.

The chassis are mounted on frames cast from LM6M alloy,* two sizes of frame being used, holding either six or four chassis. When designing the racks it was decided to fit the cathode-ray-tube storage units along the lower half of the computer, with the circuits directly associated with the storage system directly behind them, and the rest of the electronic circuits mounted on the upper half of the rack.

As the upper half of the computer could thus be narrower than the lower half, the change in width provided a platform on which the horizontal pulse-lead ducts could be mounted. With this system it was possible to arrange for each group of four storage units to be associated with a small frame holding four chassis, whilst the upper frames were larger, holding six chassis. To simplify construction the complete equipment was divided into six racks, five of which were identical, whilst the sixth was modified to deal with the special requirements of the magnetic-drum storage system.




Fig. 2.—An outline view of one of the six equipment racks used in the computer.

Each rack is 5 ft 4 in long, 3 ft 2 in wide and 7 ft high. The positions for the cathode-ray-tube stores, the horizontal pulse leads and the swinging-door type of construction are shown.

A typical rack is shown in Fig. 2, and has provision for four large chassis frames, two small chassis frames and eight cathode-ray-tube storage units. All the connections from the chassis are via connector blocks on the frames, through flexible leads, on to the centre vertical panels. The terminal blocks on the frames are of a quick-release type, into which the terminals on the ends of the flexible leads may be clipped and then clamped securely by a locking screw. The flexible leads, which consist of 44 strands of 0.012-in diameter wire, are arranged to have a bending radius of 1 in, and it is not believed that there will be any

* LM6M alloy is an aluminium alloy suitable for sand casting and gravity and pressure die castings. Its properties and chemical composition are given in B.S. 1490: 1949.

failure in these leads, since the wire has been subjected to 10 000 flexions about this radius without a single strand failing. These flexible leads are used both for power supply and pulse-lead connections. All power supplies are individually fused to each frame, the fuse blocks—taking cartridge-type fuses—being mounted at the bottom of the lower vertical panel and at the top of the upper vertical panel. The live sides of the fuses are connected to the busbars, which are of $\frac{1}{8}$ -in \times $\frac{1}{4}$ -in copper, and are positioned along the top of the rack. In addition, a 1-in \times $\frac{1}{4}$ -in copper strip along the top of the rack serves as the main earth for the system. Two of the busbars carry 115 volts 1 600 c/s at 35 amp for the valve-filament transformers, whilst the other busbars carry the h.v. supplies, where the maximum current in any line is 15 amp. The large cross-sectional area was used to minimize the voltage drop along the h.v. rails, as it was possible that the h.v. load from any part of the computer would be dependent upon the precise nature of the computation, and any voltage drop would increase the ripple on the h.v. lines, with resulting danger of unpredictable operation. The frames are hung from the rack on pin-type hinges; hence, in conjunction with the quick-release terminal blocks, it is possible completely to assemble and wire the chassis on to the frame before installation in the main racks. The flexible pulse-lead connections, of which there are 30 for the small frames and 40 for the large frames, are terminated on the vertical panels adjacent to the horizontal ducts containing the pulse leads along the centre of the racks. The vertical panels have been designed so that each lead is a continuous length of flexible wire, having the fixed terminal adjacent to the horizontal duct on the one end and the clip terminal for the frame connection block on the other. Thus the number of soldered joints in the system is reduced to a minimum.

Each lead in the horizontal duct is a single length of No. 16 S.W.G. p.v.c.-covered solid wire, the bared ends acting as the terminals. A system of Tufnol rods and moulded spacers holds each wire rigidly in position, with a spacing of $\frac{1}{4}$ in between wires. All interconnections between the pulse-lead ducts and the chassis frames are made at the centre of the racks, using leads consisting of 14 strands of No. 40 S.W.G. p.v.c.-covered flexible wire. Insulated guide rods and spacers are fitted at the centre of the rack, so that any terminal may be readily connected to any other terminal without the need for binding or otherwise fastening the wires. Since all the terminals on both the vertical and horizontal ducts are numbered, planning the interconnection wiring is virtually a clerical operation, and no elaborate wiring diagrams are required. Using a similar system, all the terminals on the chassis are numbered, as are the terminals in the connection blocks on the frames, and simple number schedules are used for describing the wiring interconnections to be made.

In the Manchester University computer the racks have been built into two parallel bays, and the pulse interconnections between these bays are made by a unit similar in design to the horizontal ducts, which passes through a trench in the floor. A spur on this trench allows leads to be brought to the console, from which the entire operation of the computer may be controlled. This console is of straightforward design, and is fitted with two small frames which carry the chassis required for operating the cathode-ray-tube displays and the input/output mechanisms.

Light-weight cover doors are fitted to the racks for all the frames, and are also hung from pin-type hinges for easy removal. These cover doors are hinged in the opposite direction to the chassis frames, so that they do not impede access to the chassis. The horizontal ducts are completely and permanently covered on erection, except at the centre of each rack where an inspection panel is fitted, held in position by spring-type fasteners. This

removable panel is found to be a useful feature in testing and maintenance, since virtually all the important waveforms used in the computer are available at all points in the horizontal duct. The panels covering the vertical distribution panels are held in position by captive screws, and are fitted with a readily removable inspection plate over the fuse holders.

In an electronic equipment dissipating 25 kW, the problem of ventilation is much simplified if the equipment does not have to be made as compact as possible. It will have been observed, from the preceding description and Figures, that a considerable amount of space has been unused throughout the equipment. Furthermore, the vertical type of construction effectively places the valves and the components in two separate "chimneys," and a reasonable amount of ventilation may therefore be obtained directly from convected air passing through the equipment. It was decided, however, that for maximum reliability an effort should be made to keep the equipment as cool as is practically possible. For this purpose a ventilation system capable of handling 3 000 ft³ of air per min was installed, and this equipment was made to serve both as an air-conditioning system for the computer room and as a cooling system for the computer.

The racks at the base, as shown in Fig. 2, were left completely open, and air ducting was fitted to the top of the racks over the busbars. This ducting was connected directly to the extraction fans. The inlet fans, to which were fitted filters and thermostatically controlled heaters, discharged clean air at a reasonable working temperature into the computer room, this air then being extracted through the computer. In an attempt to ensure that only clean air was passed through the computer, the computer room was slightly pressurized by arranging the inlet fans to pass 3 500 ft³ of air per min into it, whilst only 3 000 ft³ of air per min was extracted through the computer. It has been found, in practice, that this system suffers from three main disadvantages:

- (a) All the air passing through the computer has to be cleaned, after which it is used only once, and therefore frequent replacement of the filter elements is required.
- (b) The inlet air temperature, instead of being as low as possible, has to be held at a temperature suitable for the operators, i.e. in the range 65°–70° F.
- (c) Because of the computer and other obstacles in the room, turbulences are set up which allow air to be drawn in through open doors and windows, in spite of the excess of air being brought into the room by the ventilation system.

To obviate these faults, later computers of the same design have used a completely enclosed recirculation system, with a refrigerating plant. This enables the room temperature to be controlled independently of the computer air temperature, which may now be held down to between 40 and 50° F, and obviates the need for complex filtering systems.

(3) ELECTRONIC CIRCUIT DESIGN

The electronic circuits of the computer may be split into four main parts:

- (a) The generation of timing waveforms of basic or digit frequency, and subdivision therefrom into timing waveforms for lines, beats and bars.
- (b) Circuits used for the staticizing of control information and the decoding system used in the computer.
- (c) Circuits associated with the storage system, in particular, time-bases, store amplifiers and gating circuits.
- (d) Circuits for performing arithmetic operations.⁷

(3.1) General Design Techniques

Before giving details of some typical circuits employed, it is appropriate to mention some of the general principles employed in designing circuits used throughout the computer. The valve types used are almost exclusively the EF50 and EF55 pentodes

and EA50 diode. The EF50 is employed mostly as a squaring valve for waveforms of line or lower frequency, and as a cathode-follower in positions where either capacitive loading is not great or slow operation is of no disadvantage. The EF 55 is used as a waveform "squarer" at digit frequency, and as a cathode-follower where the lower current capacity of an EF 50 would be inadequate. The diodes are used throughout all circuits as "and" and "or" gates as well as clipping diodes. The pentodes, when operated as squarers, have only two recognized states, i.e. "on" and "off." The anode loads are so chosen that the load line on the I_a/V_a characteristics of the valve cuts the grid-voltage curves at a point well below the knee of the $V_g = 0$ curve. Indeed, it is so chosen that the anode will be at a constant potential for all permissible values of grid current, which, for the EF 50, implies a steady anode potential for grid potentials varying from -1.5 volts upwards. A pentode, so operated, will not exhibit any anode-voltage fluctuations, although the grid may fluctuate within these limits. Such a valve is described as being "bottomed." This "bottomed" condition is regarded as one state, whilst the other state is complete cut-off of all anode and screen current, which is defined to constitute not less than -15 volts applied to the grid, thus ensuring a considerable margin of safety. As has been mentioned in other Sections of the paper, the three supply potentials used in the machine are $+300$, $+200$ and -150 volts. A potential of $+50$ volts is also available, which is used mostly in digit-frequency circuits, where anode waveforms are "clipped" at this level by reverse diodes. The advantage of clipping a waveform at a potential of 50 volts over returning the anode load to a potential of 50 volts is that in either system the stray-capacitance anode-load time-constant is equal, but in the clipped case the time-constant is that for an exponential return to $+300$ volts but arrested at 50 volts. In the other case, the exponential would be of only 50 volts amplitude and would therefore produce a much slower rising edge to the waveform.

Other considerations in the design are such as to ensure that no valve or component is ever overloaded, even under failure conditions. The general design procedure is that a valve is rated at two-thirds of its anode dissipation and full screen dissipation, and is assumed to have its grid and suppressor earthed to the cathode. This may seem very cautious, especially in a valve which has, for example, a 1:1 mark/space-ratio basic waveform applied permanently to its grid. It is considered a necessary precaution, however, since in the event of a breakdown of the sine-wave generator, for example, the destruction of several valves and associated components in various positions throughout the machine cannot be risked.

The coupling between squarer valves and other circuits is usually done via a cathode-follower, the waveform being reduced to earth level at the grid of the cathode-follower. The usual method of doing this is via a voltage divider to the negative line and a catching diode to earth on the grid of the cathode-follower; alternatively, with continuous waveforms it can be done by the use of d.c. restoration. The choice of the valve and the cathode-follower load resistor is governed not only by the rated dissipation of the valve, but also by the use of the cathode-follower grid-base as a method of raising the output at least 2 or 3 volts above earth, which constitutes a safe margin in the driving of subsequent grids of squarer valves. Different parts of the machine are all d.c. interconnected, and all informative pulses are negative-going from $+2$ or $+3$ volts to at least -15 volts where the pulses are to be applied eventually to pentode grids, and going to -60 volts or more where subsequent application to an EF 50 suppressor grid is required. Many of the more important waveforms are generated in their direct and inverse forms—the inverse form being a waveform of identical timing,

but which is at a level of +2 volts when the direct version is negative, and which is negative when the direct wave is at a level of +2 volts.

Another important point in the general design procedure is the choice and rating of resistors and condensers. Ratings for resistors are halved for all dissipations, wire-wound resistors being used wherever possible for dissipations above one watt. All the carbon resistors used are of +10% ratings, whilst design is generally centred on ±20% ratings, irrespective of whether 10% carbon or 5% wire-wound resistors are to be used. Condensers are rated at 350 or 500 volts, depending on their situation, and only the highest-quality mica and paper types are used, with the exception of a very few high-permittivity ceramic types in positions where capacitance drift is of no importance. Germanium diodes have not been used, except in a few waveform-generator circuits, where their high current capacity is of advantage to their use as reverse diodes clipping waveforms in EF 55 squaring circuits.

(3.2) Squaring Circuits

Figs. 3(a) and 3(b) show the standard circuits used for EF55 and EF50 squaring circuits, and for the EF55 the use of a clipping diode is shown. Taking the case of the EF55 first [see Fig. 3(a)], the following design procedure is used. First,




Fig. 3.—Squaring-stage circuits.

- (a) The basic circuit for a squaring stage using the EF55 valve.
- (b) The basic circuit for a squaring stage using the EF50 valve.

the anode load is selected with the help of the I_a/V_a curves, in order to ensure efficient “bottoming.” This determines fully the bottom level of the waveform generated at the anode, whilst the top level is determined by the 50-volt clipping diode. This generates at the anode a waveform of 35 volts amplitude, which is independent of the main h.t. potentials. The voltage-divider chain R_1 and R_2 is used to provide an output which is at a potential above earth when the anode is at +50 volts, and is more negative than 15 volts when the anode is “bottomed.” Owing to the small excursion at the anode, it is not possible to do this when using 20% tolerances and a -150-volt supply, but as the anode voltage has been shown to be independent of main h.t. potentials, it is considered justifiable to depart from the usual tolerances in this case, and to use 10% tolerances for the resistors.

Referring to the EF50 circuit illustrated in Fig. 3(b), the anode level, when the valve is cut off, is dependent on both resistance tolerances and main h.t. supply.


Since this EF50 circuit also forms the basis for the flip-flop circuit used in the machine, with only a slight addition, and since the flip-flop circuits constitute the most complicated case, connections are shown in Figs. 4(a) and 4(b). The results of 20% tolerance calculations on the components shown in Fig. 4(a) are as follows:

Valve cut off on Suppressor Grid


- (a) Highest value of A +112 volts.
- (b) Lowest value of A +58 volts.
- (c) Highest value of B +73 volts.
- (d) Lowest value of B +6 volts.

Valve Switched On

- (e) Highest value of A +4.1 volts.
- (f) Lowest value of A +2.4 volts.
- (g) Highest value of B -68 volts.
- (h) Lowest value of B -99 volts.



(a)



(b)

Fig. 4.—Trigger-stage circuits.

- (a) The basic circuit used in the design of a trigger stage.
- (b) A trigger stage shown connected for use as a binary counter.

The flip-flop circuit is shown in Fig. 4(b). The significance of the voltage range of point B, when in the conditions (c) and (d), is that the diode connected to the suppressor grid of V_2 is bound to conduct, provided the resistances are within the tolerances calculated, and hence the suppressor grid of V_2 is bound to be at earth potential in this condition. Similarly, conditions (g) and (h) ensure that V_2 is cut off by the suppressor grid (EF50 suppressor-grid-base is approximately 60 volts).

Point A controls the level at which any trigger pulse is applied to the flip-flop circuit. On the assumption that such a pulse has an amplitude of less than 58 volts and more than 4.1 volts, it will be readily appreciated that such a trigger pulse can only be

applied to the grid of the conducting valve, even though it may have been applied, counter fashion, to both flip-flop terminals simultaneously. The condenser C_1 serves to smooth voltage changes at A and thus to ensure that it does not "move" during the application of a trigger pulse, whilst the time-constant is such that A can "move" into its new potential state before the arrival of a succeeding trigger pulse.

(3.3) The Staticizers

Staticizers are needed in the machine in order to hold control information, which is primarily in serial digit form, long enough for the orders defined by that information to be obeyed. Each staticizer corresponds to a digit position, and the presence of a pulse in that digit position is detected by "and" gating the incoming pulse train with the appropriate p -pulse. Only when an output appears on such an "and" gate is the associated staticizer triggered. The outputs of the line staticizers are between +3 and -60 volts, and the outputs of the tube and function staticizers are between -10 and +100 volts. The last two levels were chosen because these waveforms, especially with the function staticizers, have to operate a large number of gates, and the use of these levels enables the gates to be operated by resistance decoders of high impedance. The decoders are EF50 squaring valves whose grids are connected to the appropriate function-staticizer outputs via six 1-megohm resistors. Only if all the function staticizers connected to the resistors on the grid of one of these valves are at -10 volts potential can the decoder valve be in a non-conducting state. A 100-volt positive swing is used to provide an adequate safety margin for the circuit, the worst possible condition being that one resistor swinging to +100 volts potential must counteract 5 resistors of equal nominal value swinging to -10 volts potential. Under these conditions, assuming that the tolerances are adverse, the grid voltages of the valve must lie in the grid-current range in order to "bottom" it, as previously explained. A change of grid current, obtained by "moving" more resistors to +100 volts potential, will not affect the anode voltage on the valve. By using both phases of each function or tube staticizer, this switching is unique to one arrangement of the staticizers, and hence to one function code impressed on the staticizers.

(4) ELECTROSTATIC STORAGE SYSTEM

(4.1) Mechanical Construction

Cathode-ray-tube storage systems have sometimes been considered as the elements of a computer most prone to interference of mechanical, electrostatic or electromagnetic origin. With these three possible causes of disturbance in view, the present design for the stores was developed. Only electrostatic interference of comparatively great intensity is of any significance, and the only trouble encountered in the storage system has been directly attributable to the time-base generators and other waveforms applied to the store.

The storage tube is mounted horizontally in a Mumetal tube, which itself is surrounded by a tubular steel screen (Fig. 5). The two concentric cylinders are mounted on a flat plate at each end; a back-plate, which seals one end of the assembly, mounts the tube base, whilst associated components are mounted horizontally on the chassis plate. At the front of the tube there is a completely removable pick-up plate, which is spaced from the walls of the Mumetal screen by an insulating ring, which also serves as a spacer to prevent contact between the cathode-ray tube and the pick-up plate. The output is taken through the walls of the screening tubes via a screened cable, and is terminated with a coaxial plug which mates with the input socket of the pick-up amplifier. Mounted alongside this tubular assembly




Fig. 5.—A sectioned view of the cathode-ray-tube storage unit showing the mounting of the cathode-ray tube and its associated amplifier.

is the amplifier itself. This is provided with plug connections at the back, and slides into a long tray mounted on the base-plate alongside the tube assembly. The handle is designed to make withdrawal and reinsertion of the unit easy against the friction of the plug contacts at the back. To ensure immunity of this assembly against mechanical shock, the whole base-plate, complete with tube and amplifier assembly, is shock-mounted on to a cradle assembly with insulating rubber bushes, and the earth connection between the cradle and base-plate is made at one point only. The complete assembly in the cradle is mounted in a square box made of heavy-gauge mild steel, and outer covers are screwed down over the outside opening, with special provision for making good electrical contact by means of phosphor-bronze springs.

The time-bases are built into completely enclosed box chassis mounted on rubber shock mountings. Holes are cut into the top and bottom of these units and covered with fine gauze in order to provide air circulation without sacrificing electrostatic screening. The outputs from these time-base chassis are via coaxial plugs and sockets into a special duct, which runs along the whole length of the racks and serves to distribute time-base waveforms and power supplies to the stores. The duct is completely screened and time-base leads are spaced well away from one another and earth, in order to minimize stray capacitances. Connections between this duct and the stores are made via coaxial cables and plugs which mate sockets on the stores for waveforms, and multi-way plug and socket connectors for power supplies.

(4.2) Circuits in the Storage System

The circuits within the store unit, i.e. the connections to the cathode-ray tube, are completely conventional, and the amplifiers have been amply described.¹

For the X-deflection generator, a Miller valve, which has a cathode-follower in its feedback circuit, is "hesitated" during its run-down by a "dash" waveform applied to its grid, in the manner shown in Fig. 6(a). The blackout waveform applied

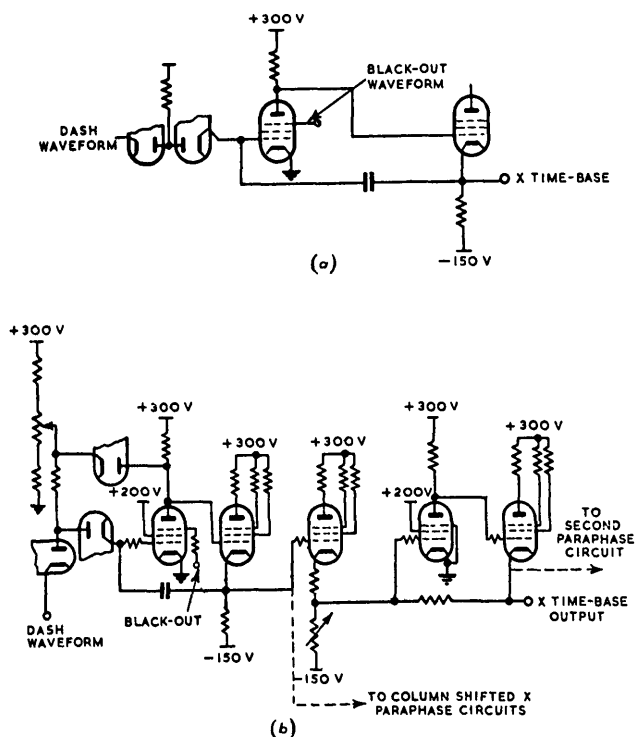


Fig. 6.—X-time-base.

- (a) The basic X-time-base generator.
 (b) The X-time-base generator, omitting the second stage of paraphasing.

to the suppressor grid of the Miller valve causes the return of the X-time-base to its positive limit during that pulse. The waveform so generated is paraphased by anode-followers, again with cathode-followers in the feedback. The precise arrangement may be seen in Fig. 6(b). With the X-time-base, two paraphasing chains are used to provide outputs from the one generator. One of these chains is used in the subsidiary stores. The other chain has a column-shift waveform injected in order to make provision for the two columns used in the main storage system.

The main Y-time-base for the stores must, of course, follow the dictates of the line staticizers during action periods and the regeneration counters during scan periods. On the input to this generator, therefore, is a system of five double-pentode switches arranged to give control of the five output lines to only one valve at a time, one of these being operative during the action waveforms and obeying signals from the line staticizers, and the other being operative in scan waveforms and obeying regeneration counters. The main Y-time-base generator is again of the anode/cathode-follower type, but the grid is connected via five diodes to a system of five resistors connected to the negative supply as shown in Fig. 7. The five controlling waveforms are applied via five more diodes to the same resistors; they are clipped so that they "move" between +15 and -15 volts. The resistors have values in ascending binary scale and are of such a magnitude that the inclusion of the largest one in the anode-follower circuit, by means of the control waveform going negative, causes the anode-follower anode to "move" some 2 or 3 volts in a positive direction, i.e. one line space. The generated Y-time-base is

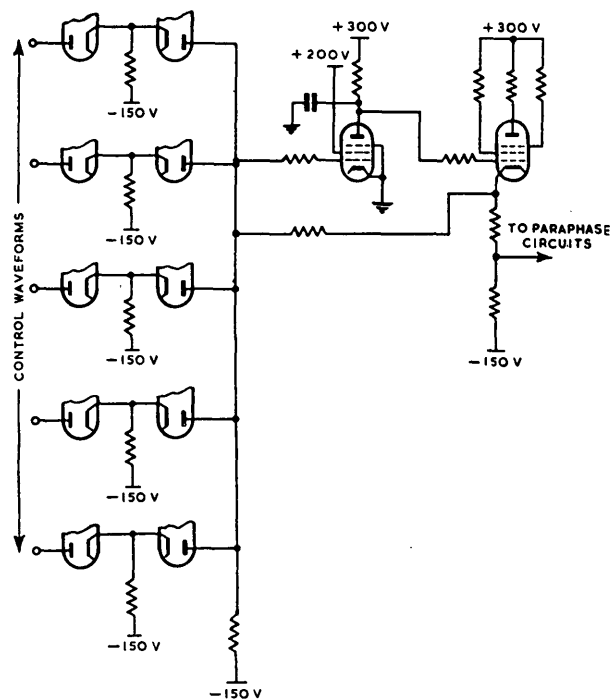


Fig. 7.—The basic Y-time-base generator.

paraphased and rephased by anode/cathode-followers in a manner similar to the single-phase X-time-base, except that the feedback resistors are of twice the value. In order to fulfil the conditions of a "dot-dash" type of storage system, a small vertical sawtooth is added into the second paraphase of this Y-time-base by feeding a suitably attenuated sawtooth waveform via a resistor into the paraphase point.

Final time-base waveforms are all centred at +125 volts, which makes the feedback resistor in the anode/cathode-followers of approximately the same value as the balancing resistor to the negative supply, which means that any disturbances of the negative supply are reproduced on the output waveform. As such disturbances cannot be tolerated, the negative supplies to the time-bases are specially stabilized by a shunt stabilizer.

(5) THE MAGNETIC-DRUM STORE

The magnetic-drum store is used in the machine to supplement the 8 electrostatic stores used for computing. Its maximum storage capacity is 665 600 binary digits, or equivalent to 512 electrostatic stores.^{5, 6}

The method of recording is by magnetizing the nickel surface of a plated copper drum rotating in synchronism with the basic waveforms produced by the machine.

The drum (see Fig. 8) consists of a copper tube 10 in. in diameter, 12½ in. in height, and with a wall thickness of ½ in. End-plate castings are fitted at either end of this tube to form a drum-like structure, and the ball races on which the drum rotates are fitted into these end castings. The central shaft, on which the drum rotates, fits the central race of these bearings and passes through the drum, being rigidly fixed vertically at either end in a heavy outer casting which contains the drum. This shaft is hollow, and on it are wound the motor and brake windings. Both these windings are wound on what would normally be regarded as rotor stampings of a small motor. The accompanying stator windings are attached one to each end-plate. These windings are of the squirrel-cage type and thus require no connection to the rotating part of the drum. Leads to the




Fig. 8.—A sectioned view of the magnetic-drum store showing the mounting of the drum on the ball races and the head mounting blocks, together with the magnetic heads and the head "write" transformers.

stationary motor and brake windings are taken through the shaft.

The recording surface of the drum is nickel-plated with great care to ensure an even and homogeneous deposit of high purity. Mounted on the outer frame, and at a distance of less than 0.0005 in from the rotating-drum surface, are dual recording-reading heads (see Fig. 9). It is necessary to maintain these heads at this close spacing from the recording surface because of the high reluctance of any air-gap in the magnetization path between head and drum. The drum must therefore be manufactured to very fine limits of eccentricity.

The drum is machined to an eccentricity of better than 0.0004 in, 0.0002 in usually being achieved. The ball races are of a specially high grade, the sphericity of the balls being 5 parts in 10⁶, and the bearing eccentricity being better than 0.0001 in. The bearings are of the angular-contact type, the bottom one sustaining the weight of the drum and the top one centring the drum, spring loading being used. Both bearings are provided with greasers, and precautions are taken to prevent the ingress of dust or grit.

In order to maintain accuracies such as these over a prolonged period, the drum must not be subjected to any sort of mechanical vibration. The outer casting, which supports it, is very rugged and provided with a heavy base-plate in order to minimize external vibrations and shocks, and the drum itself is balanced mechanically to a very high degree of accuracy. The centre of gravity is placed within 1 micron of the centre of rotation.

The combined recording and reading heads are mounted in




Fig. 9.—Details of the construction of the magnetic recording head, showing the two gaps used for reading and writing.

- A. Writing gap.
- B. Reading gap.
- C. Read-coil terminals.
- D. Write-coil terminals.
- E. Radiometal head: 0.004-in strip.

slotted strips attached to the outer casting. A spring on a shank, attached to the back of the head, exerts pressure between the head and the mounting plate, giving it a tendency to move into the drum surface. A nut which fits the threaded part of the shank projecting through the slot in the mounting-plate prevents the head from touching the drum surface and provides a fine adjustment for the gap between head and drum.

The head itself is shown magnified in Fig. 9. The writing gap and reading gap are clearly shown, as also is the single-turn coil through which writing current is passed. The reading coil is not shown; it is wound similarly to the writing coil, but has only two turns. When in position, the reading gap is in front of the writing gap, in order to provide the phase shift required in the recording system adopted.

Each of these heads writes magnetic information which occupies an annulus on the drum $\frac{1}{32}$ in wide. The spacing of the heads is determined by the slotting of the mounting strips, and is $\frac{1}{4}$ in, in order to make access to each individual head reasonably easy. Eight such mounting strips are provided, spaced at regular intervals round the periphery of the drum, each being displaced vertically by $\frac{1}{32}$ in with respect to the previous strip. The whole surface of the drum is covered in this fashion.

As an extra precaution against the ingress of foreign matter which might lodge in the gaps between drum and heads, the whole outer casting is encased in dust covers.

(6) POWER-SUPPLY EQUIPMENT

The provision of satisfactory power supplies is of the highest importance in the successful design and construction of an electronic digital computer. The specifications for such a system differ considerably from those which would be judged suitable

for, say, an analogue computer, or a complex radar system. In general, as shown in Section 3, absolute voltage stability of the h.t. supplies is relatively unimportant. The major requirement is readily seen when it is remembered that any voltage fluctuation which results in a major distortion or complete omission of a single pulse, which may be of only 1-microsec period, will almost inevitably result in an error in computation. Hence the power supplies must be completely free from rapid voltage fluctuations, whether self-generated or imposed upon the system from outside sources.

A second feature is the comparatively heavy currents which have to be handled. In the Manchester University computer, the valve heaters alone require approximately 1 300 amp at 6.3 volts, whilst the three main h.t. supplies of +300 volts, +200 volts and -150 volts have, respectively, a current drain of 10, 12 and 12 amp.

Since it would be very uneconomic to attempt to supply all the valve heaters from a central low-voltage high-current source, the use of suitable transformers throughout the equipment must be accepted. At the time when the Manchester University computer was being designed it was believed, as a result of past experience, that the cathode-ray-tube storage system was somewhat sensitive to magnetic interference from nearby transformers operated from a 50-c/s source.

The two alternatives, therefore, were to site all the heater transformers at a considerable distance from the storage units and use a 50-c/s source, or to use a higher-frequency source, when magnetic shielding would be much simplified. As the first alternative is tantamount to the provision of a single heavy-current supply, the second alternative was adopted. It was decided to use a supply at 115 volts 1 600 c/s, since suitable heavy-current generators were available. The adoption of this high frequency permitted the use of miniature sealed transformers developed for aircraft-equipment applications, and it was therefore feasible to mount a small 70-VA heater transformer on every chassis. A further advantage resulting from the provision of special generator sets was that it became a simple matter to stabilize the valve-heater supplies against fluctuations in the mains input to the motors, and also to arrange for the slow warming-up of the heaters, in an effort to reduce the stresses on the heater insulation and consequent heater-cathode insulation failure. Both these objects are readily achieved by electronic control of the exciter field current, and the system has been arranged so that, on switching on, the valve heater current is allowed to rise steadily to its normal value over a period of approximately one minute. This period is such that the peak current in the valve heaters during the warming-up period never exceeds the normal standing current; it is believed, although without any definite evidence, that this procedure has made a definite contribution to the reliability of the valves used in the computer.

It has previously been stated that the primary requirement for the h.t. supplies is that there should be complete freedom from rapid voltage fluctuations. As the current required from each of the h.t. supplies is in the range 12-15 amp, it was not thought feasible to use condensers to smooth out any fluctuations in the alternating mains supply. Furthermore, the use of mercury-vapour rectifiers was suspect, owing to their proneness to the generation of radio-frequency interference, whilst the use of hard-valve rectifiers would have been grossly inefficient. It was therefore decided that d.c. generator sets should be used, since they could be provided with sufficient inertia to overcome any mains input fluctuations, and would also supply the required current with reasonable efficiency. A further advantage lay in the simple stabilization of the generator output.

A much-simplified schematic of the power-supply system is

shown in Fig. 10. Three highly stabilized low-current supplies are incorporated in the system. These provide the reference voltages for the stabilization of the main power supplies, and are also used at one or two points in the computer, which are, of necessity, voltage-sensitive. The stabilization of the outputs from all the generator sets is achieved by electronic control of the field current; the smoothing of the output from the d.c. generator sets also uses an electronic system. In addition to the normal choke/condenser smoothing system, shunt valves with an a.c. amplifier are connected across the h.t. lines. These amplifiers have a gain of 1 050 mA/volt at 50 c/s; since the shunt valves are normally adjusted to dissipate approximately 440 mA, the system is equivalent to the provision of a capacitance of 127 mF at 50 c/s on each of the supply lines. The overall voltage stability is better than 1%, whilst the peak-to-peak ripple on the h.t. supplies is less than 0.25 volt with the computer operating.

The system is completely interlocked; it is impossible to switch on the h.t. supplies to the computer until the heater voltage is at the correct value, and the failure of any one of the supplies will result in the immediate closing down of the computer. The h.t.-supply switch, controlling the circuit-breakers, is placed on the control console of the computer, so that the supply may rapidly be disconnected from the computer in case of emergency.

Experience and developments during the last year have shown that this system, although adequate in performance, should be considerably simplified. It has been found, with improvements in the design of cathode-ray-tube storage units, that it is possible to operate 50-c/s transformers adjacent to these units without interference. The computer is now running for a minimum of 120 hours per week, and the d.c. generator sets therefore require careful maintenance if, for example, commutator ripple is not to become excessive. This system is therefore to be replaced by a static h.t. supply, using metal rectifiers, electronic smoothing and a single motor-alternator set of high inertia with a stabilized output voltage at 50 c/s. It will still be possible to warm up the valve heaters slowly, whilst the provision of a stabilized input to the h.t. rectifier units means that individual stabilization of the h.t. supplies is no longer required.

(7) MAINTENANCE OF THE COMPUTER

The usefulness of a computer is dependent upon two major factors: freedom from computational errors, which occur sporadically whilst the computer is apparently in good working order, and freedom from breakdowns which require the services of the maintenance engineer for the finding and repair of the fault.

A detailed discussion of the first factor is not within the scope of the paper. However, although the object of the design engineer is to produce a computer which is as free as possible from sporadic computing errors, occasional errors of this type will inevitably occur. The programmes for the computer must be so designed that there is an element of checking within them. This checking will increase the confidence in the final results and enable some action to be taken if a computing error occurs. In the simplest case the detection of an error will stop the computation, which must then be restarted by the operator. More refined techniques allow the computation to be restarted automatically after an error has been detected, with correction of the error, and with an indication, at the end of the computation, of the number of errors detected.

As with all complex devices, computers are inevitably prone to breakdowns, and means must be provided for the rapid rectification of these faults if a high degree of efficiency is to be obtained.

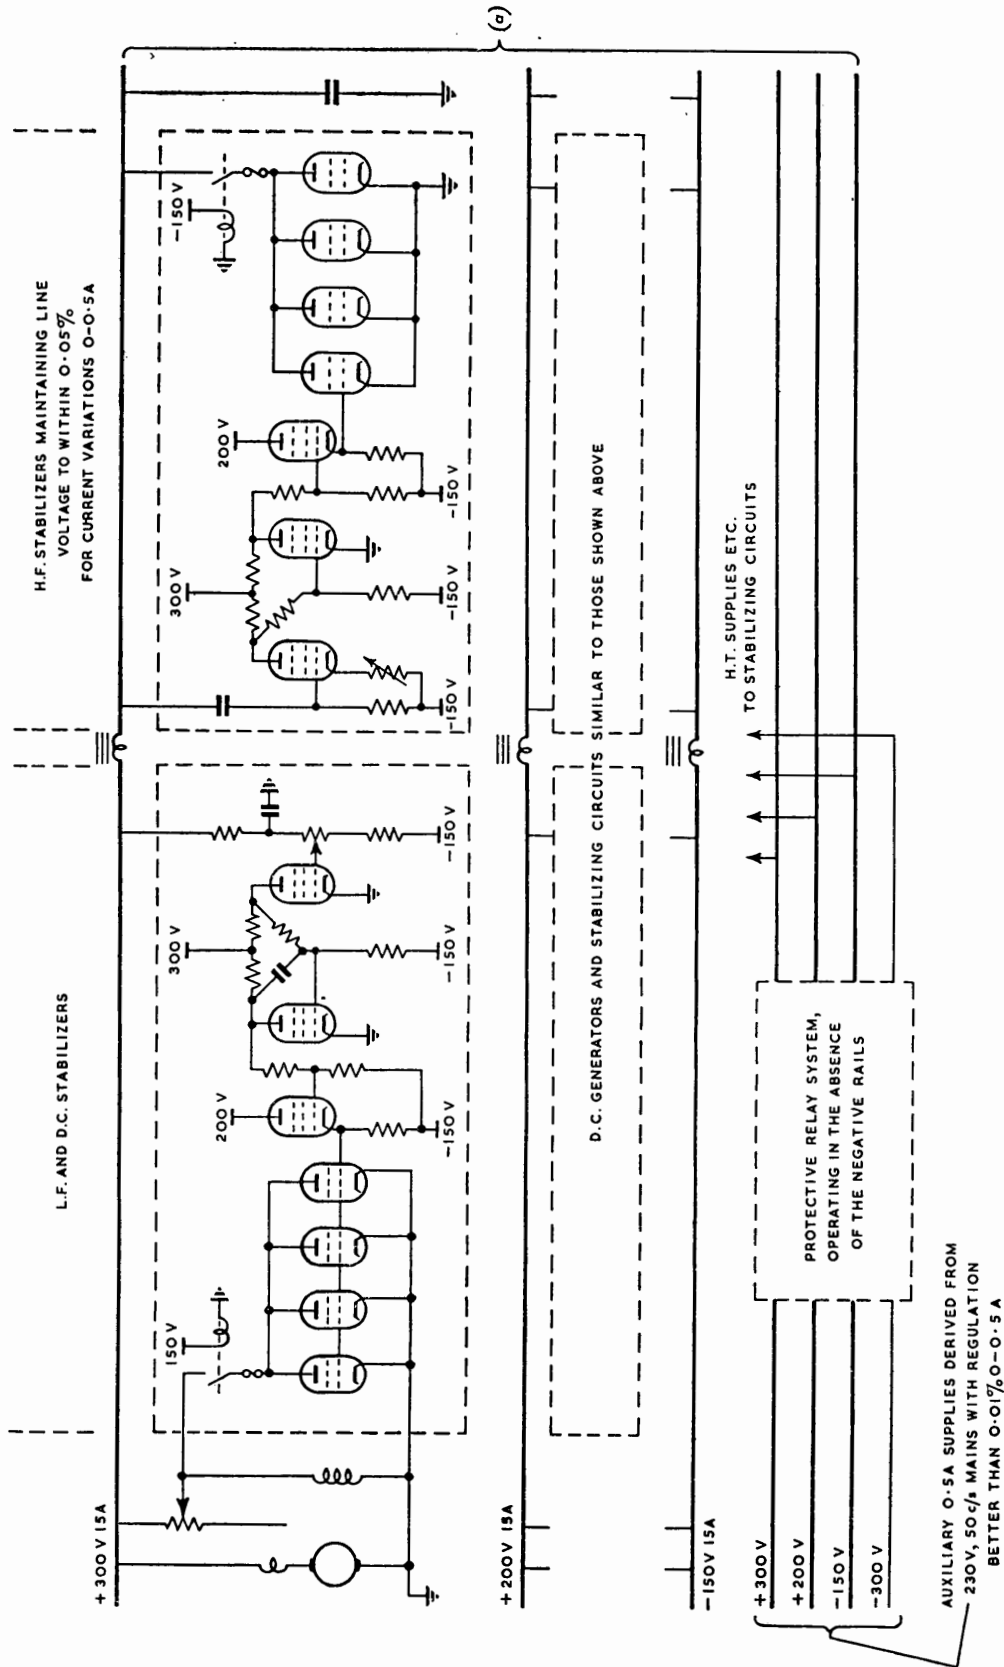


Fig. 10.—A schematic of the d.c. power-supply system.

(a) Contactors, switching lines to the computer in sequence; negative lines first on, last off. The contactors cannot be operated unless all the lines are at their nominal voltages, and all the lines are switched off automatically in the event of the failure of any one. In the event of a short-circuit on the computer side, switching off is automatically delayed for a sufficient time to rupture the line fuse in the frame where the short-circuit is located.

In general, a fault will first be apparent either as an error in computation or an inability to perform one or more of the many arithmetical, logical, or storage facilities of the computer. It is therefore the task of the maintenance engineer first to locate the source of failure, then to investigate the symptoms of the fault, and finally to rectify the fault. The location of the failure is much simplified by the use of so-called test programmes.⁸ These perform specific computations, the answers to which are known, and are so designed that any mistake in the computation will indicate, by its nature, the source of the failure. A number of these programmes are in use with the Manchester University computer, typical ones being:

(a) *Overall Test Programme.*

All sections of the computer are checked in less than 2½ min, and any major section at fault is indicated.

(b) *Store Test Programme.*

The cathode-ray-tube storage system is tested and the storage tube, or tubes, in need of adjustment, and the type of adjustment required, are defined.

(c) *Multiplier Test.*

The operation of every section of the multiplier is checked in less than 1 sec.

In addition to being of great value in locating definite faults which have occurred, these programmes are virtually the only method of locating troublesome intermittent failures, since these programmes are all designed to run continuously and may therefore be run until the intermittent failure occurs, at the moment when the faulty section of the computer is being tested.

This rapid narrowing down of the source of failure, which usually results in the fault being located to within eight valves, i.e. one chassis, and sometimes to within three valves, is not only of great help to the engineer, but also engenders confidence on the part of the operator. As all these programmes are permanently stored in the magnetic-drum storage system, they may readily be called into use, so that if an operator is having some difficulty in running a programme, he may rapidly check whether or not it is the computer which is at fault.

Once the general location of the fault has been determined, a detailed analysis must be made, usually by checking the waveforms within the circuit by means of a waveform monitor. This part of the fault-tracing procedure is, of course, made much easier by the high degree of accessibility provided by the mechanical design.

It is found, in practice, that in spite of the use of test programmes, the major part of the fault time is spent in accurately locating the fault, and it therefore appears unlikely, in a computer of this type, that the use of interchangeable, plug-in chassis, would have much effect upon the time spent in fault finding.

The performance log for the computer is analysed in the following Section, but it is convenient to trace the evolution of the present maintenance procedures. When the computer was first made available to the operators in July, 1951, it was intended to be operated during the day only. The daily procedure was therefore to switch on at 8.30 a.m., to test the computer so far as possible (the test programmes had not been fully developed at this time) and then to hand over to the operators, who worked on a fixed daily schedule of operating periods. This system was rapidly seen to have two major disadvantages, in that no time was allowed for preventive maintenance, and that an operator was not compensated for any time lost due to computer breakdown. In addition, it was difficult to arrange for normal engineering operations, such as improvements in circuit design and the installation of further equipment, without infringing upon the operator's working period. The system was therefore changed

to one in which a daily maintenance period, lasting until 10.30 a.m., was provided, and a careful check was kept upon the amount of operating time obtained by each operator. Each week's operating schedules were then based upon the results of the previous week's work. This system proved to be very effective, since the engineer was now called as soon as there was any suspicion of failure, whereas, previously, the operators had been reluctant to call upon the engineer unless the failure was catastrophic in nature.

Throughout this period it had been the practice to allow the computer to run through the evening if required, although no engineer was present. This meant, of course, that the smallest fault would result in the stoppage of work for that day. Nevertheless, the long periods of trouble-free operation during the evening, coupled with an increasing demand for computing time, suggested that greater efficiency would be obtained by running the computer on a 24-hour basis. The training of further maintenance engineers was therefore undertaken, and in February, 1952, the computer commenced to be continuously operated and maintained from 7 a.m. Monday to 10 p.m. Friday. Coupled with this, a more detailed maintenance procedure began to be evolved, which required the use of marginal test methods.

(8) PERFORMANCE AND RELIABILITY OF THE COMPUTER

Throughout the period of operation, detailed logs have been kept of the performance and of the type of faults encountered, together with the time taken to locate each fault.

Table 1 shows the performance of the computer over a period of 61 weeks. The total running time is the period for which the computer is switched on. The operating time is the period during which the computer is being used by, or is available for use by, an operator for computation purposes, whilst fault time is the time which should have been available to the operator, but which was lost due to breakdowns. Engineering time covers a number of activities. It includes the daily inspection and maintenance periods; periods which have been scheduled for the installation of new equipment, either additional to that already existing or as a replacement by a new design of existing equipment; and detailed investigation by design engineers of the operating characteristics of circuits after long periods of operation, with the object of improving the circuit-design techniques used. Hence this engineering time shows wide variations from month to month. It is inevitable that, in the first computer of a certain design to be built, a number of weaknesses will show up after it has been in use for some time, and it is essential that these should be rectified as soon as possible, so that not only this first computer, but all similar ones, will be improved as a result of the experience gained.

Table 1 also shows the number of faults which have occurred, both during the operating and the engineering time. As will be seen from Tables 2 and 3, the time spent in locating and rectifying the faults which have occurred during the operating period has been logged. It has not been practicable to do this with faults which have developed during engineering periods, since specific activities have had precedence during this time, and no immediate action has been taken to rectify faults which did not affect the tests in progress.

From Table 1 it will be seen that there have been 343 faults during an operating time of 3 791·15 hours, i.e. on average there has been a failure every 11·1 hours. Furthermore, the 343 faults have taken 553·1 hours to rectify—an average of 1·6 hours per fault.

The operating efficiency has been defined⁴ as

$$100 \frac{T_0 - T_f}{T_0} \%$$

Table 1
OVERALL COMPUTER PERFORMANCE

Period ending	Total running time	Operating time	Engineering time	Fault time	Number of faults during operating time	Number of faults during engineering time
	h	h	h	h		
1. 9. 51 (6 weeks)	341.8	215.15	114.3	12.35	9	11
29. 9. 51 (4 weeks)	264.35	202.6	37.5	24.25	33	2
27.10. 51	268.6	156.2	84.4	28.0	24	4
24.11. 51	292.6	223.5	45.4	23.7	19	2
22.12. 51	271.0	156.9	40.9	73.2	21	3
26. 1. 52	217.7	126.3	53.8	37.6	20	6
23. 2. 52	379.8	311.7	41.4	26.7	29	12
22. 3. 52	410.5	217.4	129.4	63.7	33	13
19. 4. 52	292.5	177.4	108.3	16.8	12	26
17. 5. 52	454.5	248.2	161.1	45.2	25	18
14. 6. 52	457.0	314.6	68.5	73.9	20	28
12. 7. 52	463.3	403.4	38.6	21.3	32	8
16. 8. 52	441.4	370.7	40.7	30.0	25	4
19. 9. 52 (4 weeks)	500.4	405.1	46.1	50.2	24	7
4.10. 52 (3 weeks)	341.4	262.0	53.2	26.2	17	5
Totals	5 396.85	3 791.15	1 063.4	553.1	343	149
					492	




Fig. 11.—Operating efficiency of the computer.

where T_0 is the total scheduled operating time and T_f is the fault time. Fig. 11 shows these figures plotted for the periods defined in Table 1.

Tables 2 and 3 show the types of fault which have occurred, and the time taken to rectify these specific faults. The difference in the fault time, shown in these Tables as compared with Table 1, results from disregarding, in Tables 2 and 3, the time spent in attempting to trace faults which have disappeared before the location of the fault has been determined.

It is notable that, with one or two exceptions, the average time taken to locate and rectify a fault of any type in any part of the computer has been rather less than one hour. The repair of mechanical or electromechanical faults is inevitably more lengthy than the more usual valve replacement, whilst a large change in any resistance is so infrequent that it will be the least suspected source of failure.

It will be noticed that a large number of cathode-ray-tube replacements and adjustments are shown, compared with previous reports. The large number of replacements occurred within a relatively short period, and consisted almost entirely of the replacement of cathode-ray tubes used in the subsidiary stores. A change in the manufacturing procedure for these

Table 2
FAULT ANALYSIS SHOWING TYPE OF FAILURE

Type of Failure	Total number of faults	Number of faults during operating time	Operating time lost due to faults
Open-circuit heaters	16	14	12.65
Heater-cathode short-circuit	31	27	21.25
Low emission ..	147	70	63.20
Soft valve ..	12	10	9.65
Faulty valve base ..	33	28	24.30
C.R.T. replacement ..	17	9	19.50
Store adjustment ..	74	51	47.75
Open-circuit resistor	6	5	6.25
Short-circuit resistor	1	1	0.25
Large change in resistance	7	4	8.90
Condenser failure ..	5	4	3.75
Dry joints	23	18	12.50
Preset adjustments	23	20	14.65
Mechanical or electromechanical faults	50	47	64.75
Miscellaneous faults	47	35	34.0
	492	343	344.35

storage tubes resulted in a marked tendency for the screen phosphor to lose its secondary-emission characteristics if scanned very frequently, as is the case with the subsidiary stores. This phenomenon, once recognized, was relatively simple to cure, and no further trouble from this source is being experienced. A marked increase in the number of adjustments of the storage system during operating periods was noted after the introduction of marginal testing procedure. The normal marginal test applied was the reduction of the voltage applied to the heaters throughout

Table 3

OCCURRENCE OF FAULTS IN VARIOUS SECTIONS OF THE COMPUTER

Section of computer	Approximate number of valves in section	Total number of faults	Number of faults during operating time	Operating time lost due to fault
				h
Waveform generator ..	930	68	37	41.95
Control	180	17	9	11.80
B tube _{2,4}	180	17	9	11.55
Accumulator	300	36	28	28.50
Multiplier	850	54	37	48.25
Magnetic storage system	870	58	44	33.95
C.R.T. storage system ..	350	128	91	87.90
Input/output equipment	200	72	56	41.65
Console	50	18	13	8.85
Power supplies	150	24	19	39.90
	4 070	492	343	344.35

the computer. It was found that this had a delayed effect upon the storage tubes, and steps had therefore to be taken to isolate the storage-tube heater supplies from those to the rest of the computer, which can now be varied independently.

An account of the performance of a computer would not be complete without some reference to the efficiency obtained in the computation of problems taking many hours for solution. A university computer, which is not primarily intended to supply a computing service, will normally be used on research problems, where the major part of the time will be devoted to the development of the programme, and the actual computing time will be fairly small. A number of lengthy computations have, however, been undertaken, and a measure has been obtained of the computing efficiency. In these "production runs," efficiency has

Table 4

USE FACTOR FOR LONG PRODUCTION RUNS

Month of 1952	Total production time for which figures are available	Use factor
	h	%
February	14	45
March	28	43
April	No long production runs were made, due to extensive engineering work	
May	63	60
June	50	80
July	64	97
August	27	100
September	70	93

[The discussion on the above paper will be found on page 540.]

been defined as the ratio of the theoretical time that the computation should take, to the actual time taken, no allowance being made for time lost due to normal maintenance, mistakes by the operator in handling the computer, computer faults, or mistakes made in computation, i.e. a perfect computer has been assumed as the yardstick. Table 4 shows some results obtained during the last six months, together with the actual length of each production run.

(9) CONCLUSION AND ACKNOWLEDGMENT

Details have been given of the more important design features of the universal digital computer at Manchester University, together with the operational results which have been obtained during the working period.

The steady improvement in performance is best shown by the production-run figures, and it is felt that figures of 80% or greater, obtained from scheduled production periods, are a good indication of the generally satisfactory design of this computer.

The authors wish to acknowledge their indebtedness to Prof. F. C. Williams, and the computer research team at Manchester University, for advice and encouragement, and to the Ferranti computer section, especially Mr. G. I. Thomas, Mr. B. G. Welby, Mr. H. Malbon, Mr. E. T. Warburton, Mr. F. Cooper, Mr. G. Fox and Dr. A. A. Robinson, for their many contributions to the design of this machine. They also wish to thank Ferranti Ltd., for permission to publish the paper.

(10) REFERENCES

- (1) WILLIAMS, F. C., and KILBURN, T.: "A Storage System for Use with Binary-Digital Computing Machines," *Proceedings I.E.E.*, 1949, **96**, Part II, p. 183.
- (2) WILLIAMS, F. C., KILBURN, T., and TOOTILL, G. C.: "Universal High-Speed Digital Computers: A Small-Scale Experimental Machine," *ibid.*, 1951, **98**, Part II, p. 13.
- (3) WILLIAMS, F. C., and KILBURN, T.: "The University of Manchester Computing Machine," Review of Electronic Digital Computers; Joint American I.E.E. and I.R.E. Computer Conference, Dec. 10th/12th, 1951.
- (4) KILBURN, T., TOOTILL, G. C., EDWARDS, D. B. G., and POLLARD, B. W.: "Digital Computers at Manchester University." See page 487.
- (5) WILLIAMS, F. C., KILBURN, T., and THOMAS, G. E.: "Universal High-Speed Digital Computers: A Magnetic Store," *Proceedings I.E.E.*, 1952, **99**, Part II, p. 94.
- (6) WILLIAMS, F. C., and WEST, J. C.: "The Position Synchronization of a Rotating Drum," *ibid.*, 1951, **98**, Part II, p. 29.
- (7) WILLIAMS, F. C., ROBINSON, A. A., and KILBURN, T.: "Universal High-Speed Digital Computers: Serial Computing Circuits," *ibid.*, 1952, **99**, Part II, p. 107.
- (8) GRIMSDALE, R. L.: "Computing Machines: Design of Test Programmes," M.Sc. Thesis 1951, University of Manchester.