The Making of Colossus

ALLEN W. M. COOMBS

The article describes the techniques that evolved for the design and manufacture of the Colossus machines and their auxiliaries, invented for the purpose of code breaking during World War II and installed at Bletchley Park.

*Categories and Subject Descriptors: K.2 [**History of Computing**]—Colossus, hardware, people*
General Terms: Design
Additional Key Words and Phrases: cryptology, Bletchley Park, "Heath Robinson"

1. Introduction

In the first of these three papers, T. H. Flowers has described how he was by a series of lucky chances brought into the field of code breaking, and how his prewar experience and his ideas on the subject of electronic switching were thereby enabled to burgeon in the making of what were in some sense the first electronic computers in the world. In the third paper, W. W. Chandler will expound the problems of installing and commissioning the machines at Bletchley Park, against the clock, employing construction units made straight from the drawing board, and with staff who had to learn their expertise as they went along. This second paper links the other two; I hope to show how the team set out to design and build large machines of a totally novel form in minimal time, inspired simply by the firm belief that this could be done, that it was the right thing to do, and that by guess or by God we were the people to do it.

A point is first to be made concerning historical sequence and timing. I was enrolled in the Colossus team in or about September 1943, having until then been engaged in frustrating the knavish tricks of His Majesty's enemies (as we say in our national anthem) in other and rather less interesting ways. At that time the first Robinson machine was waiting to be made to work prior to installation, and the first (and unofficial) Colossus, the Mark I, had been completely designed and was approaching manufacturing completion in the laboratory at Dollis Hill (see Flowers's paper). It was to be my duty (I now gather) to concentrate on the Robinsons exclusively; however, the very successful demonstration of the Mark I Colossus in December 1943 changed all that. From then on my principal work was on the Mark II. I was a member of the design team, and also in charge, under Flowers, of its detailed production and manufacture. It follows, therefore, that comments in this article on the original design of either the Mark I or the Robinsons are inferences; I had no direct hand in them, but I feel sure nevertheless that my deductions are reasonably correct.

2. The Preexisting State of the Art

The Robinson designers were clearly of the opinion, usual in those days (Flowers himself would appear to have been the sole exception, an experience not unfamiliar to him), that whereas valves (vacuum tubes) had certain advantages in speed and other things, still they were fragile and sensitive devices, to be eschewed as far as possible. A colleague once said to me, "Valves? Don't like them. Nasty things. They break." The result of that attitude is typified by the "modulo-2" addition circuit used in Robinson (see Figure 1).

The switching logic in this circuit was performed by changing the phase of a 25-kilohertz supply through 180 degrees in the presence of a signal Z from a photocell in the bedstead, or leaving it unchanged in the absence of such a signal—phase modulation, in

Author's Address: Blyth Knoll, Yealmbury Hill, Yealmpton, Devon, PL8 2JX, England.
© 1983 AFIPS 0164-1239/83/030253-259$01.00/00

fact. A succession of switches driven from different photocells should produce an output in exact phase or exact counterphase with the original voltage (depending on the modulo-2 sum of the photocell signals) and, provided that each switch has a gain of unity, directly comparable with it in magnitude. In the diagram, block B is a Hartley oscillator providing the 25-kh source voltage to as many loads as may be required (hence the buffering attenuator network) at a level of about 0.5 volts. Block A is one switching stage of a number in tandem, the required connections being made by telephone cords (shown dotted). The transformers in any one A block are carefully and accurately tuned to give no phase change from input to output, except as may be specified by signal Z, which can produce a complete reversal by operation on the ring modulator shown. The gain of the valve circuit is adjusted (by means of feedback resistor R) to be exactly unity from input to output of the adder when it is properly terminated. The voltage resulting from several stages of switching is then compared with the original supply in the comparator; in the case shown, it is arranged that if the modulo-2 sum of the input signals is zero (an even number of phase reversals), a null output is obtained, whereas a sum of one (an odd number of phase reversals) gives reinforcement of the supply, and hence a rectified signal at the output. The Robinsons worked at 2000 characters per second, and the photocell output was for about two-thirds of a signal pitch; at 25 kh there were thus about nine cycles of signal ideally available for detection.

Allen W. M. Coombs received B.Sc. and Ph.D. degrees from Glasgow University in 1932 and 1936. He is an associate of the Royal College of Science and Technology (now the University of Strathclyde). From 1936 to 1973 he was employed by the British Post Office (now British Telecom) on various projects in its Research Branch. He spent the last 10 years or so of his employment in research on intelligent machines, in connection with pattern recognition for postal mechanization.

Coombs writes the following about the present: "He is now blissfully retired to the glorious county of Devonshire, where he and his wife can grow their own food, make their own wine, sail their own boat, and get thoroughly engulfed in the maelstrom of English village life, and whence they can make periodic excursions in foreign travel. And despite the recrudescence of many wounds from his student days, he still believes that the Almighty created heaven and earth with Rugby Union football in mind."

The principle behind this circuit would seem to have been that each stage of modulo-2 addition could be obtained by the use of a single valve together with a number of other less variable though still quite accurately determined components such as resistors, transformers, rectifiers, and capacitors. The switching logic was not performed by the valves at all, but by the rectifiers of the ring modulator, the function of the valves being merely that of providing unity gain per stage. In theory, the circuit was very ingenious; in practice, it was a nightmare. Each stage had to be most carefully adjusted to give exactly correct phase and gain (involving endless experiments with cathode resistors), and even then, when they had been set up as accurately as possible, by the time a number of stages had been put in tandem the output was neither in phase nor in counterphase, but somewhere in between. The effect was very strange; about six stages in tandem would work perfectly, but the addition of two more—any two—would produce a meaningless phase change. So the whole set of adjustments had to be done again, with precisely the same result. This particular problem was solved (though that may not be quite the right word) by a piece of inspired serendipity on the part of Flowers. I remember going to him in despair after a fortnight of unavailing effort to produce circuits that would work either singly, in sixes, or in more than sixes. His instant solution was, "Change the frequency." I did. It worked. I still don't know why. Nor, he tells me, does he.

Even when the circuits were working properly, there was still a major problem arising out of the components used. It has been shown that ideally about nine cycles of 25-kh voltage were available. But the circuits contained inductances and capacitances as well as nonlinear components, and they were switched with step-function voltages; they were therefore prolific transient producers. Nor was it certain that the signals from the photocell tape readers would be exactly synchronous (in fact, they weren't), yet within nine cycles it was necessary that the output signal stabilize sufficiently to give a firm difference between something and nothing for a strobe pulse to operate on—and the whole thing had to restore to zero before the next signal started.

We achieved it—mostly by inspired juggling with resistors. But it was quite clear that noisy and edgy circuits like this would not do for the higher signal frequency (5000 characters per second) to be used in Colossus. Flowers's ideas culminated in the Mark II circuits, which had the following facilities.

1. Isolation between circuits, when required, provided by buffer valves, not by resistor networks. This

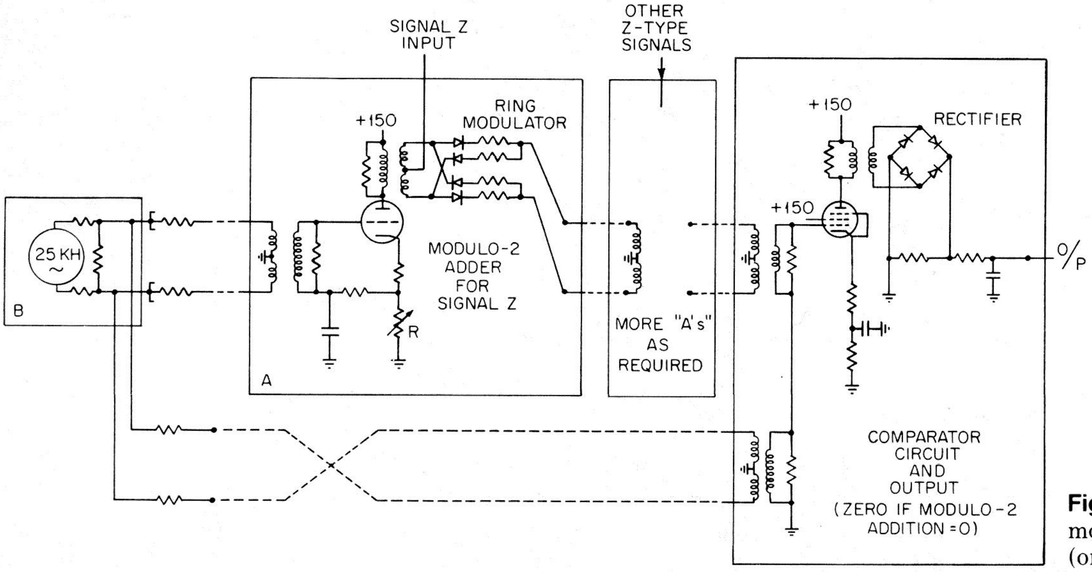

Figure 1. Robinson modulo-2 adder (one valve per input).

clearly implied a readiness to multiply the number of valves, contrary to the established practice of the time.

2. No inductances in critical circuits, though electromagnetic relays in slow-speed circuits were not precluded. With only resistors and capacitors, transients are simple to calculate and control.

3. No representation of signals by bursts of oscillatory voltage. Instead, signals were transformed as soon as possible after the bedstead stage into direct-voltage swings of considerable magnitude, and extended into what later became known as "non-return-to-zero" state.

4. No exact adjustments of gain. The resistors available to use were in any case not free of drift in value—nor the valves of variation in emission—so signals were at all stages derived from valves which were fully "bottomed" (that is, with their anodes down as far as they could be driven by running into grid current) or else fully cut off (that is, with no anode current at all). The first of these two states is achieved most satisfactorily with pentode valves (in which anode current is dependent within limits more on screen voltage than on anode voltage), which inclined us to prefer pentodes and enabled us to use the last facility.

5. Use of the suppressor grid as a switching means (see next section). These ideas, together with the substitution of one tape by thyratron rings (see Flowers's paper), solved the Robinson difficulties, but they naturally introduced some of their own. What the difficulties were, and how we dealt with them, will be treated next.

3. The Designing of Colossus

The major advantage of transformers is that the absolute voltage level of a signal becomes unimportant; our adoption of exclusively direct-current signals meant that we were immediately faced with the fact that a signal swinging about earth at one stage of switching became (from the nature of valve circuitry) a signal swinging about a higher voltage (about +70 volts in our case) at the next. If several stages of switching were involved, then either the absolute level of voltage had continually to rise or the level had somehow to be dropped (by direct-current means) between stages. In later years, some machines did indeed use the variable-level technique, to the extent of having voltage rails ranging over 1200 volts or more; we deemed ourselves short enough of manpower anyway, and we limited ourselves to voltage rails between −150 and +200 volts. This meant that the voltage-dropping requirement was critical, for a change in level by resistive potentiometer means also a drop in voltage swing of the signal, and the available swing was not always big enough to be reduced in this way, the more so since the circuits had to be secure against the initial tolerance and drift aging characteristic of the commercial resistors available to us under wartime conditions and in bulk.

We devised a method of dealing with this problem, as will be shown. At this point, however, it is appropriate to indicate another of our design standards arising largely out of the resistor limitation. Reduction

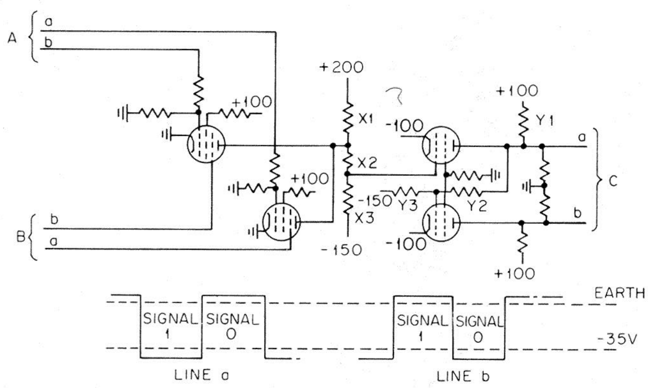

Figure 2. Modulo-2 adder and signal form.

of voltage swing due to a single potentiometer network is troublesome enough; reduction in a multiple network (that is, one in which several sources are connected through resistors to a single output) is clearly far worse, and we avoided such circuits in all except one case, that of the "scale-of-five" counter. Instead, we isolated the signals by what came to be called the "one-electrode, one-job" technique. If two signals were to be combined in some logical way, they were fed to two separate valves with a common cathode or anode. Better still, they were fed to one valve, but separately to the control and suppressor grids, for we had realized that a pentode valve could be cut off on either, thereby giving a logical "and" for upward-going signals and a logical "or" for downward-going signals. The suppressor technique was so useful, and we became so attached to it, that after the war we tended to avoid the double triodes used by most computer makers for their binary stores, on the grounds that such "long-tailed pairs" left no separate electrode for switching purposes (apart from other disadvantages). It has to be admitted, however, that the suppressor swing to achieve cutoff was much greater than the control-grid swing, and the circuits had to be arranged to cope with this.

The result of our deliberations was a set of standard conditions as follows.

1. Interstage signals to be carried by pairs of wires, one wire at slightly higher than earth potential, the other at less than −35 volts. The binary signals zero and one were represented by reversal of the voltage conditions on the two wires.

2. Signals to be fed out from buffer valves of low impedance, thus allowing multiple feed-out, cutting out "back circuits," and avoiding the worst effects of wiring capacitance (and therefore delay) on interrack wiring.

3. Normal switching valves to be small pentodes (code EF36), with triodes for cathode followers and heavier-duty valves for buffered outputs.

4. Heater voltage to be available at two voltage levels, namely, earth and −100 volts. It was permissible to run the heater and the cathode of any one valve at different voltage levels provided that the heater was not positive relative to the cathode, and provided that the difference did not exceed about 60 volts.

The standard screen voltage for a conducting valve was about +80 volts relative to the cathode, derived through a dropping resistor from +100 volts. This gave a control-grid base (that is, cutoff voltage) of about −6 volts and a suppressor base of about −35 volts, with an anode-current capability of at least 3 milliamperes when the anode was fully bottomed (at about 15 volts). Control grids were run into grid current in the conducting condition (which in practice meant that they stood at −0.4 volts relative to the cathodes).

The next section indicates the effect of these design criteria on some typical circuits.

4. Some Typical Circuits

4.1 The Modulo-2 Adder

It is appropriate to start with the modulo-2 adder (Figure 2) since it gives an interesting contrast to the Robinson adder (Figure 1). The essential circuit is that given in Flowers's paper, with a pair of output valves added to satisfy condition 1 in section 3. Two potentiometer chains, the X's and the Y's, are involved.

4.2 The Shift Register

One stage of shifting, or "remembering one back," is shown in Figure 3. The idea is that a signal A on the input pair of lines is to be read off at a given instant by the strobe pulse B (derived from the sprocket output) and recorded on the bistable D (output C), while at the same time signal A is itself possibly undergoing change under the influence of the same strobe pulse B—hence the delay networks E, which prevent the new signal at A running right through the circuit to the output. Observe that potentiometer X is somewhat more complex in this case, requiring tighter control on the component selection and tolerance, a factor now to be considered.

In the two cases cited, it was found that satisfactory voltage swings could be obtained with resistors varying ±10 percent from nominal. But most of the resistors available to us in 1943 (at least in quantity) were guaranteed only to ±20 percent, and were in any case

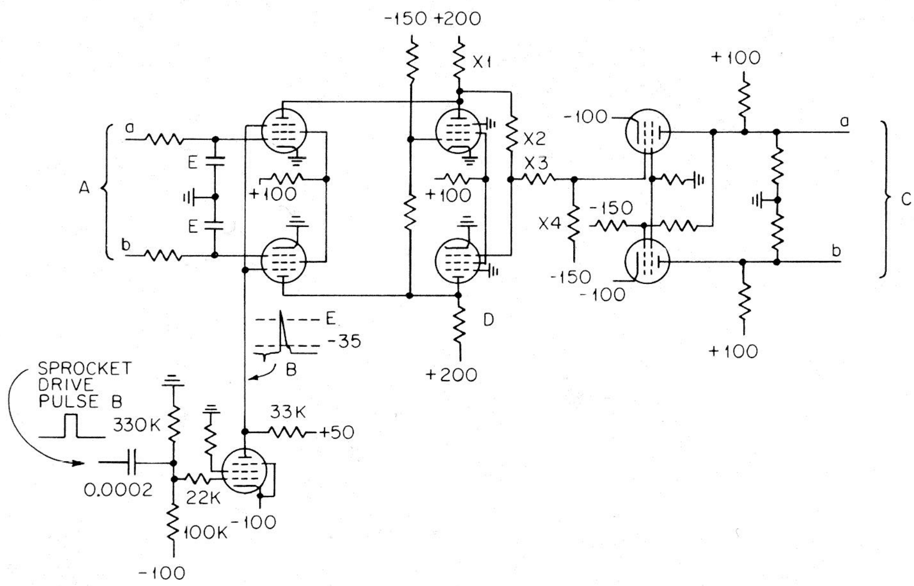

Figure 3. Shift register.

known to be subject to 5 percent drift with age (meaning within six months). Since we had to make the machines at maximum speed using large numbers of resistors, but without sacrificing reliability, this looked like an impasse. However, we found that the makers, while not able to work to tighter tolerances, were nevertheless prepared to send us resistors already sorted into 5 percent groups, though we would have to accept the groups as they came off the production lines. Thus resistor No. 28 was nominally 220 kilohms; we would receive supplies of it in eight marked boxes, box 28a containing resistors of value from 180 to 190 kilohms, and so on up to box 28j, which contained resistors of value from 250 to 260 kilohms. With the potentiometers designed to nominal values, the rule was that the resistors in any one machine were to be taken from boxes with the same suffix letter (or, failing that, resistors on one rack, or, failing that, on one plate); choosing the boxes to be as near nominal as were available, we could be reasonably sure that the resistor ratios were always provided by ±2.5 percent resistors, and that even with a subsequent drift of 5 percent the voltages derived from the potentiometers would never exceed tolerance. In fact, they never did.

There was one exception to the rule just cited, which has already been mentioned: the scale-of-five counter. Here we had to depart from the other rule (against multiple potentiometers), and we found that the tolerance in the resistors feeding the valve grids had to be much tighter. A description of this circuit follows.

4.3 The Scale-of-Five Counter

The ring of five valves constituting the scale-of-five counter (Figure 4) was a pentastable (though I don't remember that we ever called it that). The equilibrium conditions were with any one valve cut off and all the rest bottomed. It is a major problem to calculate the available grid voltage swings under all possible adverse conditions of resistor variation with such complex circuits, and really the only absolutely safe tolerance for the resistors of the multiple potentiometers was ±1 percent—with no drift. A few such resistors were available, but not nearly enough (we required 500 per machine), so we had to make do with 4 percent ranges (±2 percent) of the best quality we could get. It helped that we used lower-valued resistors than usual in the screen circuits (10 kilohms for 22 kilohms), which had the effect of bottoming the anodes to 12 volts, and, as it happened, the counters were satisfactory; we were probably saved by statistical averaging of the variations in each resistor network.

Observe the use of suppressor-grid switching in this circuit. The counter-drive pulse (of about 10 microseconds duration) was applied to all the control grids in parallel and had the effect of driving those grids positive and the valves into grid current. Four of the valves were of course already in this condition, however; the remaining one—the one marking the present state of the counter—responded by bottoming on its anode, thereby producing a differentiated negative-going pulse at the suppressor of the next valve in the ring, which cut that valve off. The suppressor pulse

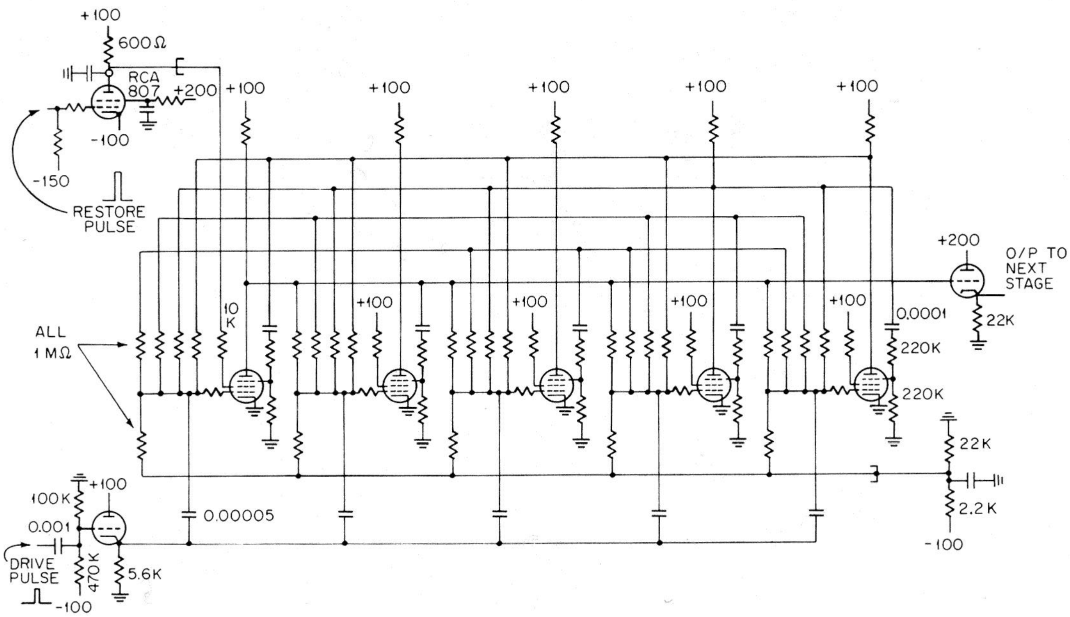

Figure 4. Scale-of-five counter stage.

was timed to outlast the drive pulse (40 versus 10 microseconds), so that when the pulse drive had been removed and equilibrium reestablished, the counter had moved on by one step only.

We may have been luckier than we realized in this circuit. In postwar times we tried to use valves with short suppressor-grid bases, to spare us the large voltage swings on such electrodes found necessary in Colossus. Valves of this sort by that time existed; however, whereas the EF36 was insensitive to the first 8 volts or so of negative swing on the suppressor, the postwar valves were not. The effect was that sudden swings of the anode were reflected by capacitative feedback into the high-impedance suppressor circuit, and the valves could go into sustained multivibration. But we had no such problems with Colossus.

4.4 The "Dropper" Circuit

The output from the bedsteads was no more than 0.5 volts, at earth level. This raised the only really difficult direct-current amplification problem in the machine, for the signal had to be raised to standard magnitude, still at earth level, and bottoming the amplifier valve was not feasible at such a low input. Our solution was to use that property of the pentode valve which preserves a fairly constant anode current over a wide range of anode voltage. Figure 5 shows valve A having a nearly constant screen voltage and grid bias, and therefore producing a fixed voltage drop in the anode load, largely independent of the anode voltage; as a result, the anode voltage changes of amplifier valve B are reproduced with little loss, but at a lower level, at the grid of signal-shaping valve C. It was necessary to adjust bias resistor D of valve A in order to place the grid voltage at the right level (valve-emission variation being mostly responsible for this), but the adjustment was not unduly sensitive, since the swing was by this time large, and I do not think it gave any trouble, once made.

5. The Manufacturing of Colossus

The work was planned on a sort of "critical-path-analysis" basis (though we didn't call it that then), with the parts likely to take the longest time started first—even before we received official authority to proceed, as Flowers has said. The circuits were designed by Flowers, Chandler, Sidney W. Broadhurst, and myself—I remember Flowers having made a rough draft of the projected new machine, then tearing his draft into pieces and sharing out the bits to us with the instruction to get on with it. Once completed, the designs were handed over to other engineers whose function it was to lay out the circuits for assembly on standardized plates. The assembly and wiring were carried out by technical staff, including some of the wartime female assistants, first at Dollis Hill and later also at the Birmingham factory, with one of our engineers (O. G. T. Belcher) acting as resident supervisor. Racks were wired by small squads under more experienced skilled workmen, and then the whole lot was sent for assembly and commissioning at Bletchley.

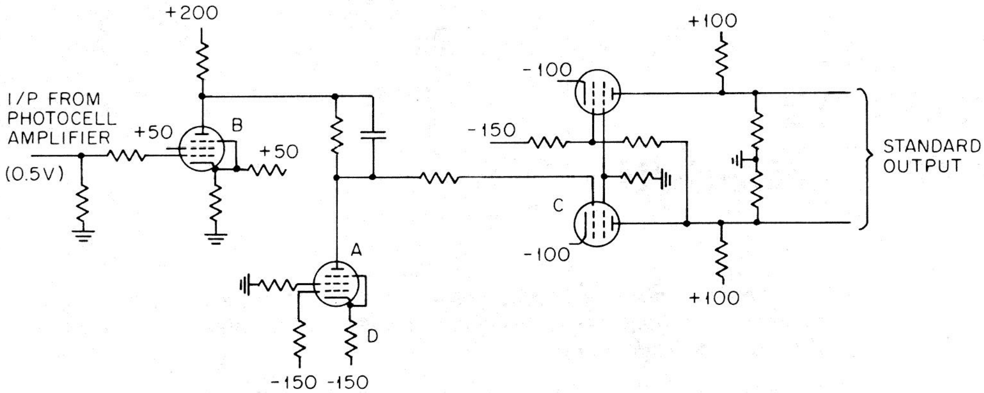

Figure 5. Pentode "dropper" circuit.

About 70 people of many ranks were involved, of whom some 15 finally went to work at Bletchley as maintenance staff. All these spent long hours on work of such a nature that its very purpose could not be revealed to most of them (one bright spark with a knowledge of radio carefully drew out the circuits he had been wiring up on a plate, and came to the conclusion that he had been making a lot of direct-current amplifiers, which of course he had), and from June 1944 they had to contend with the extra hazards of flying bombs and V2 rockets. But the work never stopped; the job was done and the machines delivered, which reflects great credit on those concerned, especially (when you come to think of it) on those who didn't even know what they were doing, and to that extent might naturally have lacked motivation.

There were, of course, other circuits using less novel forms of switching to control the slower-speed parts of the machine. Here again luck was with us. We were not only the Post Office Research Station, knowledgeable about teleprinters; we were also the Signalling Group, with a wealth of experience on automatic telephone exchanges and such devices as electromagnetic relays and rotary switches, and we had at least one genuine wizard in the person of Broadhurst. Both he and Flowers were the sort of men who could at the drop of a hat design highly complicated relay circuits to perform any operation required—and get it right first time.

We were in the best Shakespearean tradition a few, a happy few, a band of brothers. I had at that time just had the great happiness of being accepted by my future bride (and present wife); to grace the occasion, Broadhurst solemnly named one of his relays VJB (her initials), this being a relay having a strong interaction with another named AWMC. *Eheu, fugaces, Postume, Postume, labuntur anni.* [*"Ah, Postumus, Postumus, how the fleeting years go by."*—Ed.]

6. And Lastly

So the designs were completed, and such was our confidence in the new technique that we rapidly reached the stage where, new requirements having been formulated or new ideas advanced, we would specify forthwith the new circuits required in the sure and certain knowledge that they would work. Of such a nature was the special machine we called Proteus, another called Aquarius (which stored information as charges in small capacitors we continually had to top up—whence the name), Salamander (which was never completed, having been forestalled by the peace), and the attachment called Mighty Wurlitzer, which can be seen as a sort of cinema organ keyboard near the foreground Wren in Figure 14 of Flowers's paper. The habit of confidence was hard to break. After the war, as a member of a totally different team on totally different work, I once blandly suggested that such and such would be the circuit I would use, and that would be all right, wouldn't it? I mean, they wouldn't want a mock-up or anything, would they? I was speedily disabused. War is one thing; peace is quite another.

So nothing remained but to assemble the machines and make them work. But that was the department of Bill Chandler, whose paper follows.