The Design of Colossus

THOMAS H. FLOWERS

During World War II the German armed forces used machine-enciphered teleprinter messages for some of their high-level communications. Mathematicians at Bletchley Park, a high-security establishment in Britain, discovered processes by which such messages might be decoded; for the information to be useful, however, processing speeds such as only electronics could attain would be necessary. The first electronic machines made for the purpose did valuable work but were too slow and cumbersome to handle all the traffic being received. They were superseded by faster and more versatile machines called Colossus.

This article describes the construction and operation of the Colossus machines. The machines had most if not all of the essential features of a modern computer, except that variable programming was provided not by memory store but by hard-wired function units selected and interconnected by switches operated by the mathematician-programmers.

*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"

Foreword

Howard Campaigne

My view of Colossus was that of cryptanalyst-programmer. I told the machine to make certain calculations and counts, and after studying the results, told it to do another job. It did not remember the previous result, nor could it have acted upon it if it did. Colossus and I alternated in an interaction that sometimes achieved an analysis of an unusual German cipher system, called "Geheimschreiber" by the Germans, and "Fish" by the cryptanalysts.

For its time Colossus was a notable innovation, different from its predecessors in many dramatic aspects.

1. It was electronic. Other machines, such as Heath Robinson, had had electronic subsystems, but they were of minor size. Colossus had 2400 vacuum tubes—big bottles. Ah, the warmth at two A.M. on a damp, cold English winter!

Author's Address: 12 Sunnyfield, Mill Hill, London NW7 4RG, England.

Illustrations for this article are courtesy of British Telecom Research Laboratories. The article is adapted from a lecture given at the Digital Computer Museum in Marlboro, Mass., on October 15, 1981.
Acknowledgment is made to the Director of Research of British Telecom for permission to make use of the information in these articles.
© 1983 AFIPS 0164-1239/83/030239–252$01.00/00

During World War II, Howard Campaigne was with a U.S. Navy communications operation and was assigned to work at Bletchley Park. After the war he joined the newly formed National Security Agency, where he eventually became chief of research. He left NSA in 1970 to become professor and chairman of the Mathematics Department at Slippery Rock State College in Pennsylvania. He retired in 1976 and now lives in Portales, New Mexico.

2. It was digital, and experience with digital circuits was then very limited. The vacuum tubes of the day were mainly intended as amplifiers; manufacturers strove for linear response. Fortunately for Colossus they were successful over only a limited range.

3. It was programmable by means of a switchboard. Toggle switches enabled one to choose among binary functions of the input, which was a long string of cipher text, and then the outputs of these functions were counted. At the end of each pass of the input string, the counts were used to control the printer, suppressing those counts of lesser interest.

Colossus was not sequence controlled, nor was its program stored. Nor did any of us see that possibility.

Parameters of the cipher were entered by means of bottle plugs on the remote back of Colossus—to keep leads short. The cipher text was on a very long teletype punched tape pasted into a loop and then run at 5000 characters a second. When a paper tape parted it caused some excitement.

The tapes were prepared and mounted by women of the Women's Royal Naval Service (WRNS)—called "Wrens"—who were very skilled and adroit. I tried but invariably caused trouble. Without the Wrens I was helpless.

Colossus was useful in more than one way, and there were even demonstrations applying it to number theory. But these demonstrations were more notable for their ingenuity than for their effectiveness.

A successful result from Colossus was not plaintext, but an intermediate product that was completed by hand by skillful specialists.

The cipher system under attack was on-line, an integral part of the communications links; a typist at one end ran a typewriter at the other. Once synchrony had been established, the typist fed in message after message until the backlog had been exhausted. A cryptanalytic solution would reveal an avalanche of plaintext.

This cryptanalysis was a superb technical achievement, and the cryptanalysts were very proud of it. At the same time they were painfully aware that their dominance was precarious, and were fearful that a change by the Germans might deprive them of their sustenance. When I was scheduled by the U.S. Navy to join the "Newmanry" as a working member and observer, my plans were unsettled by a dispatch that the Germans had begun to use new wheel patterns each day instead of once a month, and that there was little chance that the system would ever be read again. But two months later the effort had been doubled and redoubled, and more was being read than ever before. Part of the redoubling was to build more Colossi; earlier there had been one, now there were to be twelve.

F. H. Hinsley says in "British Intelligence in the Second World War" (Vol. 2, p. 29), "In terms of the intelligence value of their contents . . . the decrypts of Geheimschreiber were commonly more important than the Enigma; moreover a Fish decrypt commonly incorporated a large number of individual signals. Whereas the bulk of the Enigma traffic was at and below Army level, the Fish links were confined to communications at Army level and above, that is between Armies and Army Groups and OKH; this being so, they carried orders, appreciations and situation reports, and even routine returns of strengths and supplies, of which the decrypts were of exceptional significance."

Howard Campaigne
1809 South Main Street
Portales, NM 88130

During World War II, I became involved in codebreaking activities for which I conceived and built machines that became known as Colossus. No public mention of these machines was made at the time. It was 30 years later that Brian Randell, professor of computing science at the University of Newcastle upon Tyne, chanced to hear about them. All the record of their existence then available was in the memories of those who were involved at the time and still surviving. Randell was able to meet and question many of these people and thus gained enough information to piece together a description of the construction and operation of these machines. Realizing that the machines were in fact electronic data processors, he raised the question of their relevance to modern computers.

That Colossus occurred at all was the result of a series of lucky chances to which one more is added if because of it I am to be numbered among the computer pioneers. At the time I had no thought or knowledge of computers in the modern sense and had never heard the term used except to describe somebody who did calculations on a desk machine.

Prewar Work

The real origin of the Colossus machines was the work I did before the war at the Dollis Hill Research Station of the Post Office, which ran Britain's telephone and other communication services. Soon after I went to the laboratories in 1930 my researches were directed to finding the reasons why so many of the calls dialed on the automatic exchanges that were then spreading rapidly throughout the country gave wrong connections or no connections. I think my reaction to this work was that of any young engineer of the time: how can we make a better system?

By 1935 I had become interested in electronic switches as an alternative to the electromagnetic switches then universal. This was not so remarkable as it might seem. Copper oxide rectifiers and thermionic vacuum tubes had become plentiful and cheap in the 1920s and by 1935 were being used as switches in a variety of applications. In speech and radio transmission systems they were used in modulators for their high switching speed, and in telegraph transmission systems as relays because of their reliability, which is what I was looking for. High-speed counters using tubes had been invented for scientific experiments, the first in 1929 by C. E. Wynn-Williams at Cambridge University. As far back as 1919 the Eccles-Jordan trigger circuit had been published but with practically no use having been made of it even though it could be used as a relay and as a 1-bit store. (Perhaps I should say here that when I use modern terms such as *one-bit store, program, programmer*, and so forth, I am not implying that they were in use at the time I am talking about, but in the hope that they will make what I am saying easier for a modern reader to understand.)

From 1935 onward I was exploring the uses of electronics in telephone exchanges. By 1939 I felt able to prove what up to then I could only suspect: that an electronic equivalent could be made of any electromechanical switching or data-processing machine. I well remember the elation I felt when I reached that point because it meant that not just parts of a telephone exchange could be electronic—complete exchanges could. Something more than equivalence was clearly going to be needed, however, because exchanges based on that principle were much too costly to be practical. I knew nothing of similar ideas being pursued elsewhere because none had reached the stage of publication, but it now seems that none were further ahead than I was.

Enigma

Hostilities commenced in Europe in 1939, and all civil work in Britain had to be subordinated to war work. In 1942 I was sent to Bletchley Park, a highly secret establishment about 50 miles north of London, to take on some top-secret work. By good luck I was about to be involved in code breaking and would encounter a problem for which my prewar work would be needed no matter what the cost.

From about a year or so after the war had started, cryptanalysts at Bletchley Park had been reading German messages that had been encoded on a machine called Enigma. Alan M. Turing was a leading figure in this activity, and he explained it to me and told me what the Post Office was required to do—something that presented no difficulties to a telephone engineer. The work was soon found not to be needed; had that been realized earlier (and there was some suspicion that it could have been), Dollis Hill might never have been involved with Bletchley Park. Although the final outcome of the war would no doubt have been the same, its history might have been different with greater loss of life and damage. Two other Dollis Hill engineers, Sidney W. Broadhurst and William W. Chandler, and I had become involved, however, and we were available when a new engineering problem arose.

The Germans had invented a machine to encipher teleprinter messages using a principle invented by an American, Gilbert Vernam, in 1917 during World War I that had scarcely been used because of operational difficulties. The German machine looked like and was operated like a teleprinter but differed in that messages sent in plain language from the keyboard were transmitted by the machine in cipher, and enciphered messages received in the machine came out printed in plain language, all automatically without effort on the part of the machine operator other than setting the machine to some initial mechanical conditions specified in a codebook before sending or receiving a message. Operationally nothing could have been simpler.

Bletchley Park received the messages on 5-hole punched tape, as shown in Figure 1. Each character is a 5-binary digit number represented by holes in lines across the tape together with a small sprocket hole by which the tape is driven so that it may be punched and subsequently read. Additional holes were put by Bletchley Park into the blank ends of the tape preceding and following the message—a start hole between the third and fourth digits and a stop hole between the fourth and fifth digits to indicate to the processor when the message was about to start and when it

Thomas H. Flowers trained as a mechanical engineer before entering the Post Office Communications Laboratories in 1930. There his main interests were first telephone direct-distance dialing and later electronic exchanges. He continued this work with ITT-England from 1964 until retiring in 1970. He received a B.Sc. from London University in 1933 and an honorary D.Sc. in 1977 from the University of Newcastle upon Tyne for his work on Colossus.

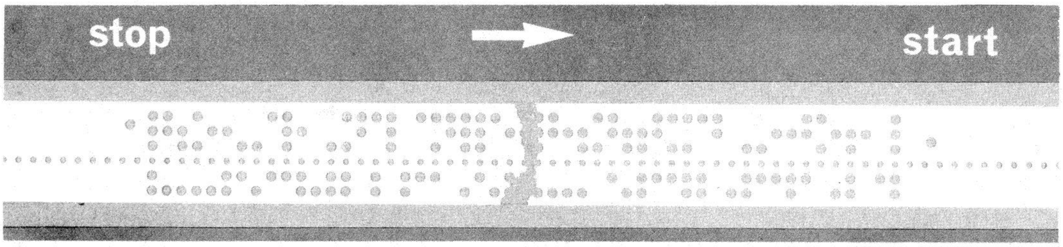

Figure 1. Message tape with start (right) and stop (left) holes.

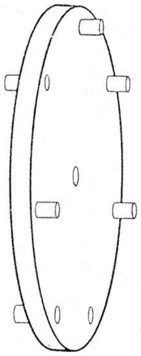

Figure 3. Code wheel.

ended. A punch specifically made for the purpose is one of the few surviving relics of the Colossus machines (Figure 2).

Enciphering and deciphering within the machine depended on 12 code wheels around the peripheries of which were two-position cams that could be set by hand to one of the two positions (Figure 3). By these means each wheel stored a pattern in binary code, the numbers of digits in the patterns being all different because the wheels were of different diameters, and all prime to one another to maximize the number of different relative positions of the wheels. As the wheels were rotated during the transmission of a message, the cams operated contacts to vary the five digits of the teleprinter code in a pseudo-random sort of way. Each of the five digits was enciphered individually, as shown in the diagram for one of the digits (Figure 4). The plaintext value of the digit for a character sent without encipherment is A, which in binary notation may be 0 or 1. The value of the contacts B operated by the cams of one of the wheels may similarly be 0 or 1, depending on the position of the cam effective at that moment. The output from the B contacts is the modulo-2 addition A + B to which another addition is made by contacts C controlled by a cam on another wheel, the output from the contacts C being the enciphered digit to be transmitted. The electronic equivalent of the relays is also shown.

For the five digits of the teleprinter code, there were two sets of five wheels operating cams B and C. One set of wheels took one rotational step for each character transmitted. The second set stepped not for every character but erratically as determined by the cams on two further wheels. The number of combinations of different wheel positions and therefore of characters that could be transmitted before the cipher repeated itself was so great that in practice it never did repeat. The enemy had such confidence in the invulnerability of messages enciphered in this way that they used these machines to send information of the highest importance such as battle orders in great detail and even messages and instructions from Hitler to his field commanders.

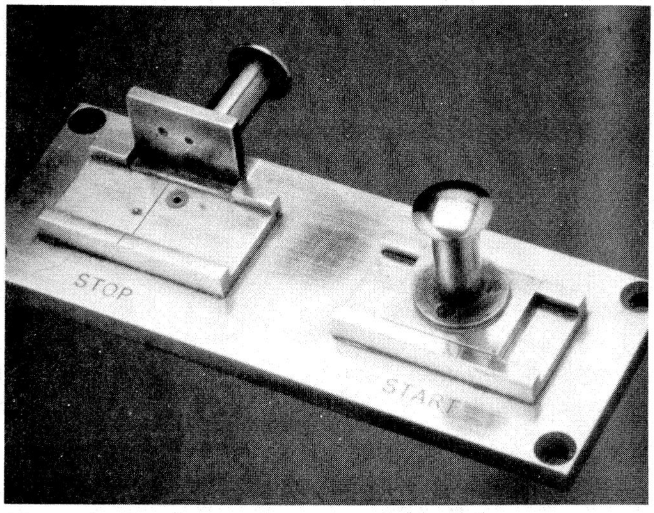

Figure 2. Punch to make start and stop holes.

Electronics

Maxwell H. A. Newman, who at one time had had Turing as a mathematics student at Cambridge University and after the war became professor of mathematics at Manchester University, was brought into Bletchley Park late in 1942 to work on the German teleprinter traffic. The problem was to find first the wheel patterns then in operation, at that time being changed periodically, and next the starting positions of the wheels for every message received. Newman had the idea that a "key" tape reproducing some characteristic of the coding machine and formed into a loop could be processed with a message tape also formed into a loop (Figure 5); if the processing were repeated as many times as were necessary for every possible position of the "key" tape relative to the message tape to be tested, some clue might be obtained. Synchronization between the two tapes would be maintained by driving them by two sprocket wheels mounted on the same shaft and engaging the sprocket

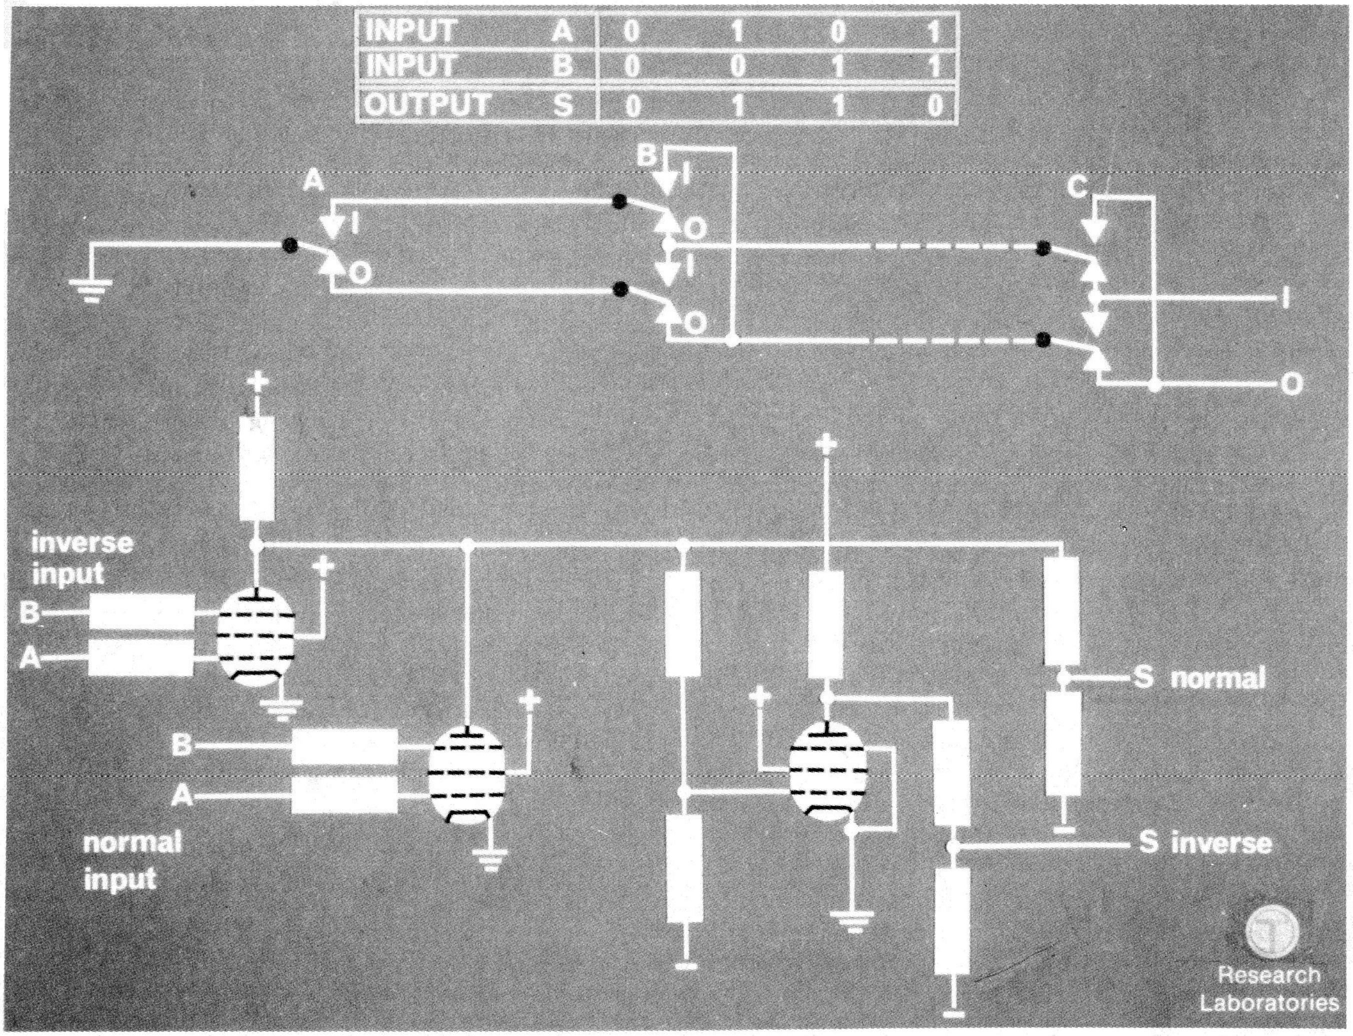

INPUT	A	0	1	0	1
INPUT	B	0	0	1	1
OUTPUT	S	0	1	1	0

Figure 4. Modulo-2 addition.

holes in the tapes. Every relative position of the two tapes would occur in time if the numbers of characters in the two tapes were made prime to each other. With one clue and another "key" tape, another clue might emerge. Continuing in this way, enough clues would lead to a solution.

The quantity of processing was manifestly exceedingly large: hence even to test Newman's theories, some kind of machine had to be devised. It was easily calculable that a machine using the fastest techniques of data processing available at that time would take too long to produce operationally useful results. Electronics was clearly necessary, and Wynn-Williams, the inventor of the first electronic counter, was brought into the problem. F. O. Morrell from Dollis Hill, an expert in teleprinter engineering, was put in charge of the project. The loops of tape were too long to hang freely, so a structure with movable pulleys was needed to carry them and was called a bedstead.

Photoelectric reading of the tapes was the only possibility at the speed at which they had to run. Two Dollis Hill physicists, Eric Speight and Arnold Lynch, designed, and the model shop made, a photoelectric reader that was similar to a cinematograph projector but projected in the opposite direction, from a lamp and optical system mounted on the gate to a screen inside the machine that was a photographic plate opaque except for some clear spaces behind which photocells were located. One of these screens, which were really masks to control the light reaching the photocells (Figure 6), is still available although broken. The crescent-shaped clear spaces that occur in the mask were designed so that in passing over the mask, circular areas of light derived from holes in the tape generated approximately rectangular output pulses from the photocells and amplifiers. Ten photocells in two rows of five enabled two successive five-unit characters on the tape to be read simultaneously,

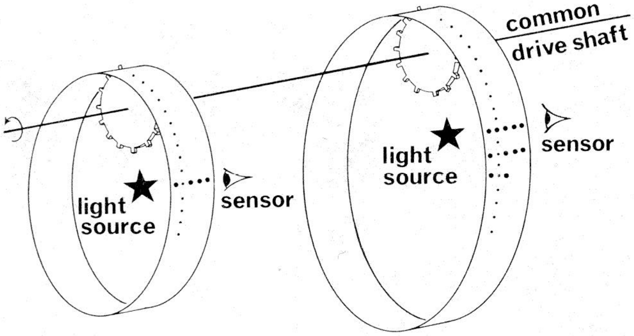

Figure 5. Loops of tape run in synchronism.

a feature required of some of the processing. The sprocket holes produced a continuous output of pulses from another cell and amplifier, and the start and stop holes put into the tape were detected by two further cells and amplifiers in paths between the second and third and third and fourth digits.

Driving the tapes by sprocket wheels was tried experimentally as soon as a bedstead with tape-driving gear became available. To give some idea of the problems involved, typical tapes would be 5000 and 1000 characters long, for which one complete program of processing would necessitate 1000 revolutions of the first simultaneously with 5000 revolutions of the second. It seemed even without actually trying it that the chance of paper or any other kind of tape standing up to passing over the sprocket wheels that number of times without damage was remote; the initial tests were in fact distinctly discouraging. Various designs of sprocket wheel were tried with no conspicuous improvement.

With diminishing prospects of ultimate success, some new ideas began to be urgently sought. At this point I realized that my prewar work pointed to a solution to the synchronization problem that was simple in principle although it might be difficult to put into practice. It was not the paper tape that was essential to the processing, but the data on the tape. All that the "key" tape was doing was storing pregenerated data so that they could be read into an electronic processor at a much higher speed. An electronic equivalent of the machine that had originally generated the data—and I had no doubt that such an equivalent could be made—could work fast enough to issue the data directly to the processor and thus render the "key" tape unnecessary. Only one loop, that of the message tape, would be needed, and that could be driven by smooth-faced pulleys at high speed without fear of damage. Moreover, the considerable trouble, time, and labor needed to make "key" tapes and load them into the tape reader would disappear.

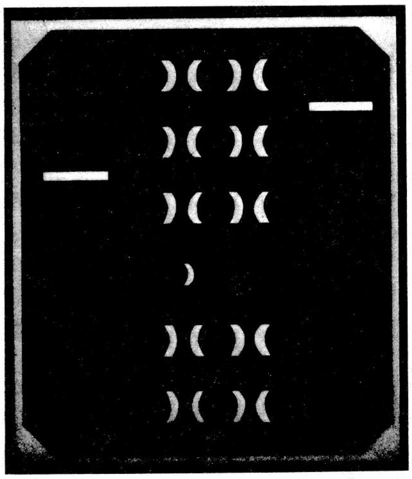

Figure 6. Mask in photoelectric scanner.

My suggestion, made in February 1943, was received with considerable scepticism. The first reaction was that a machine with the number of tubes that was obviously going to be needed would be too unreliable to be useful. Fortunately, this criticism could be defeated by the experience of the Post Office using thousands of tubes in its communication network. These tubes were not subject to movement or handling, and the power was never switched off. Under these conditions tube failures were very rare.

Some people suggested that the war would be over before the first machine could be operational, but the Post Office engineers thought they could make such a machine in one year, and it was unlikely that the war would be over by February 1944. Success could not be guaranteed for a machine for which there was no precedent, however. Hence it was decreed that work on the two-tape machine should continue and have first priority. What to do about a one-tape machine with all-electronic processing would be left to those who would have to make it to decide.

Events justified this decision. Machines with two tapes operating in synchronism at 2000 characters per second were in service operation from about the middle of 1943 and were code-named after Heath Robinson, a cartoonist well known for his drawings of ludicrous tasks performed by fantastic machines. The tape-running problem had been solved by motoring one of the smooth pulleys to relieve the sprockets of most of the driving power but still able to maintain synchronism. Much valuable work was done by these machines toward validating and developing the ideas of Newman and other mathematicians recruited for the project. The quantity of messages decoded was small in relation to the total traffic but nevertheless important.

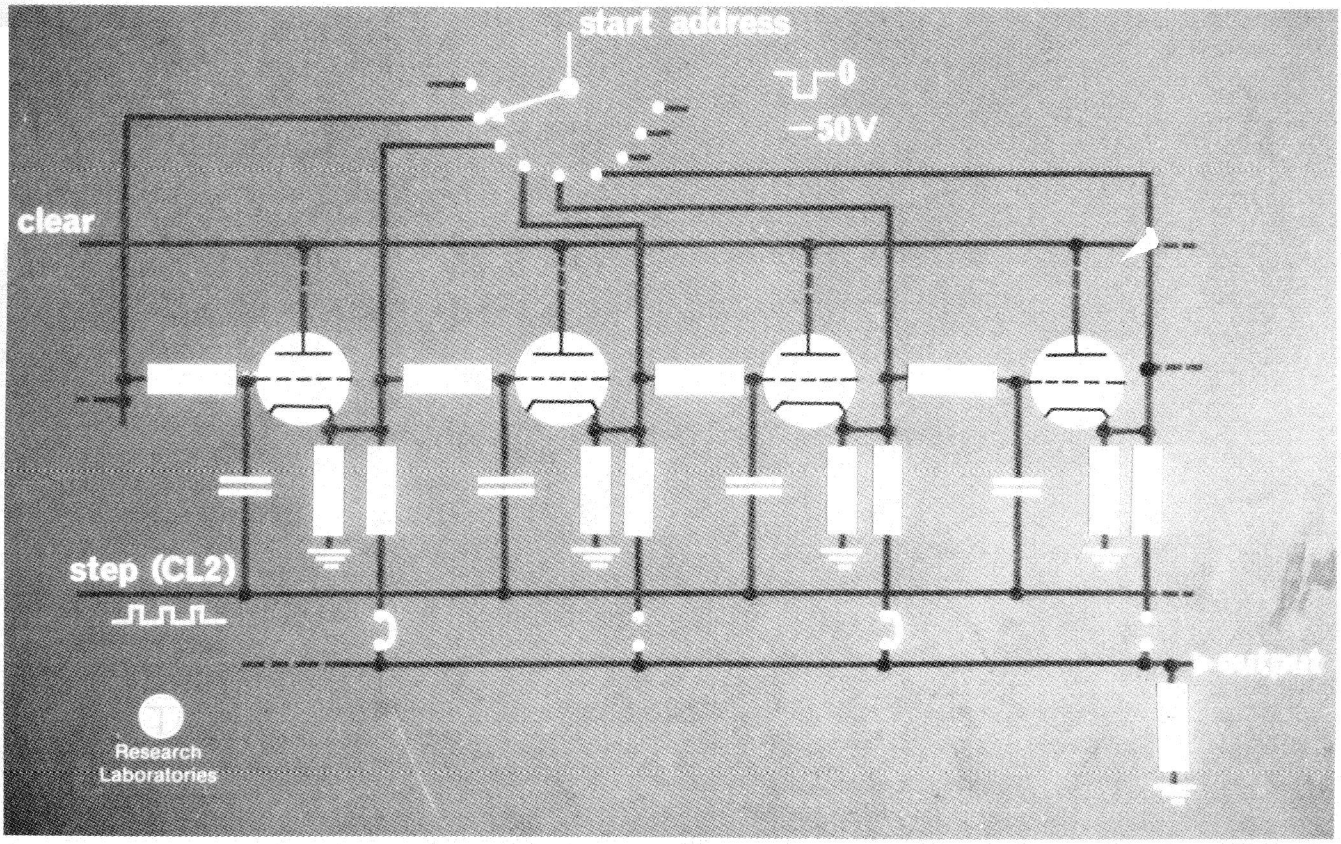

Figure 7. Portion of thyratron ring store.

Meanwhile, the Dollis Hill engineers, having decided to make an all-electronic processor, took only 11 months to get the first machine into service, a feat made possible by the absolute priority they were given to command materials and services and the prodigious efforts of the laboratory staff, many of whom did nothing but work, eat, and sleep for weeks and months on end except for one half day per week that they allowed themselves for the necessities of life such as talking to their wives and getting their hair cut. The United States also contributed tubes and an electric typewriter under the lend-lease arrangements.

Colossus

Responsibility for the overall and detailed design of the first machine was shared between Broadhurst, Chandler, and myself, with Allan W. M. Coombs joining us for subsequent design and construction work. The first machine went into service in the early part of December 1943 and was code-named Colossus by Bletchley Park because it was the largest machine they had had to operate until then. The cryptanalysts were impressed by the speed of operation, which was 5000 characters per second, and by the reliability, which soon gave them confidence in the results produced. As the cryptanalysts became more familiar with the new machine and its greatly increased processing power compared with the Robinsons, ideas for new processes occurred, some of them being impossible to execute on the Robinson machines.

The practical speed at which Colossus could be operated was tested by loading it with a maximum-length tape and increasing the speed until something happened. At about 9700 characters per second, the tape broke in one and then several places. The various sections of tape did their best to obey Newton's first law and travel in a straight line in the direction they happened to be going with an initial velocity of nearly 60 miles per hour, and thus found all sorts of curious places in which to come to rest. Clearly, the maximum safe speed at which paper tape could be driven could not be much greater than 5000 characters per second. A much higher speed had been aimed at to reduce the numbers of machines eventually to be built, and the operating force to practical quantities, thus prompting the thought that the effective speed could be increased by parallel processing. By using five processors in parallel, all operating on the same program but with different input data from the tape, an effective speed

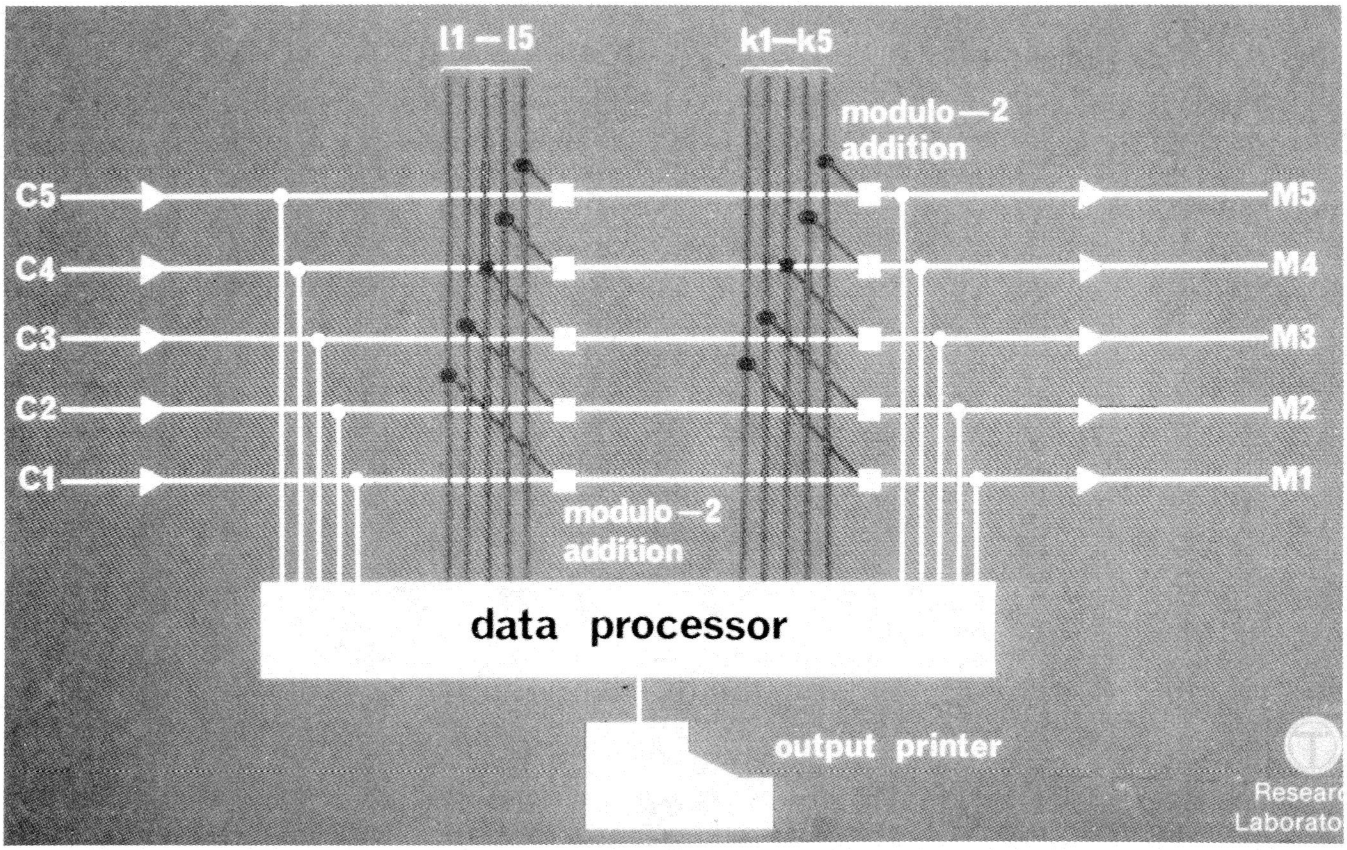

Figure 8. Simplified schematic of Colossus.

of 25,000 characters per second was obtained and was enough. Five separate processors presented little difficulty. A photoelectric reader to supply them with data would have to read six lines of tape simultaneously, however, and would be very difficult to construct. Six lines were needed so that two consecutive but different lines could be presented to each of the five processors. This difficulty was solved by shift registers invented for the purpose. Data read from the tape were read into six-bit shift registers, which meant that six consecutive characters from the message tape were available for processing using a tape reader that had to read only one line at a time and was thus simplified compared with the readers first used with the Robinson machines.

We anticipated that several more Colossus machines would be demanded in a great hurry if the first was successful. For that reason a redesign incorporating improvements and new ideas began even before the first machine was finished, as did the manufacture of parts, particularly those parts that took a long time to make. At about the end of February 1944, at a meeting at Bletchley Park, I was told that 12 machines were to be in operation by June 1. The instructions had come from the highest level, and I well remember the consternation when I had to say flatly that it was impossible. This sudden demand had stemmed not only from the success of the existing Colossus machine, but also from increasingly widespread German use of teleprinted messages, and from "improved" German operating procedures, which meant that we had to work harder to achieve the same results. Because some parts were already made, we could promise that the laboratories would make one or two machines and have the first in service by the required date. For the remainder, the Post Office factory in Birmingham would have to be organized to make most of the parts for assembly at Dollis Hill, and delivery at the rate of one per month after the first was working was thought to be possible.

After another period of intensive work by all concerned, the first of the new machines was operational on May 31—except that it was subject to intermittent and mysterious faults. Just after midnight of that day, members of the commissioning team were too tired to go on and went home for a few hours of sleep. When they returned at 8:30 the next morning, June 1, they found the machine in service. What had happened was that a Dollis Hill engineer on the night shift had discovered the trouble to be parasitic oscillation be-

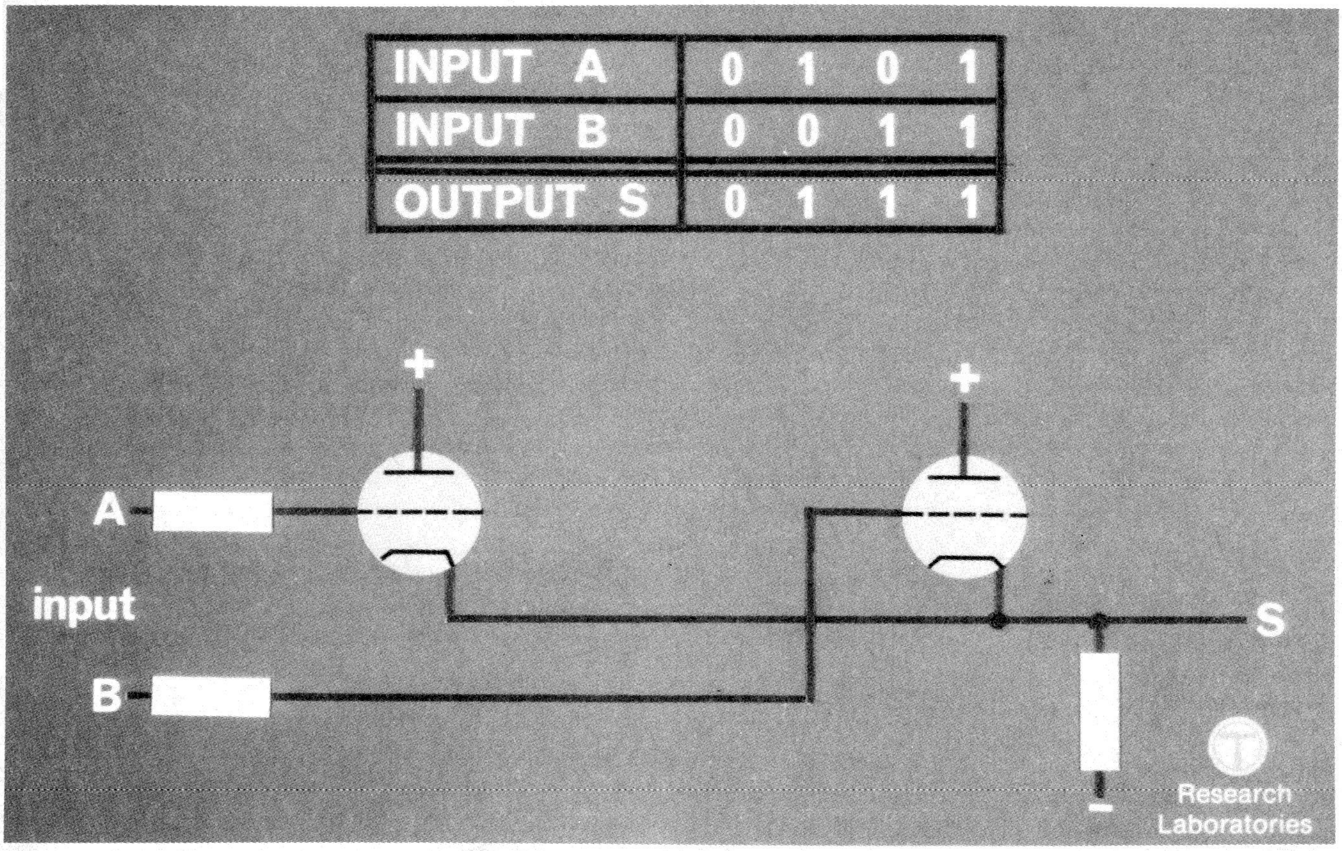

INPUT A	0	1	0	1
INPUT B	0	0	1	1
OUTPUT S	0	1	1	1

Figure 9. Boolean addition.

tween the heaters and the cathodes of some of the tubes and had cleared the trouble by the insertion of a few resistors.

I have included these details of what happened in the hope that they convey some of the dedication, the excitement, and the job satisfaction of those hectic days. The remainder of the machines did in fact follow as promised, so that when the German surrender 11 months after June 1, 1944, put an abrupt end to the whole enterprise, 10 machines were in operation and the 11th was being commissioned.

A principal element of Colossus was the bit-stream generator, which replaced the Robinson "key" tapes. It comprised a bank of 12 rings of different scales, which generated the bit streams from semipermanent data. The bit-stream generator and the data processing were necessarily all electronic to attain the processing speed required, but electromagnetic telephone-exchange-type switches were also used in large quantities for slower-speed operations simply because that was the quickest and best way of making the machines. The first Colossus had contained 1500 tubes; the production machines needed 2400 tubes because of their increased processing power and facilities, together with 12 large rotary switches, about 800 relays, and an IBM electric typewriter to print the output. The semipermanent data were read in by U-links inserted by hand into pairs of sockets and read out by thyratron tubes connected in series to form a ring operating as a Wynn-Williams type of counter (Figure 7). Thyratron tubes have a grid electrode by which they can be fired—that is, caused to conduct current from anode to cathode—and once fired they stay fired until the power supply to the tubes is cut off.

Each thyratron was connected to one of the pairs of sockets such that when the tube was fired, it caused current to flow to the output of the ring if the socket contained a U-link. Only one tube could be in the fired condition at any one time, and as pulses to be counted were received, the fired tube moved step by step around the ring. At each step, one bit of stored data was read to the output. In operation during processing, when one revolution of the message tape and scan of the message had been completed, a clear pulse applied to the anodes of all the thyratrons of all the rings ensured that none was in the fired condition. A negative pulse applied through the start-address switches of the rings to be involved in the processing caused one thyratron in each to be fired. These rings were then ready to count sprocket-hole pulses occurring

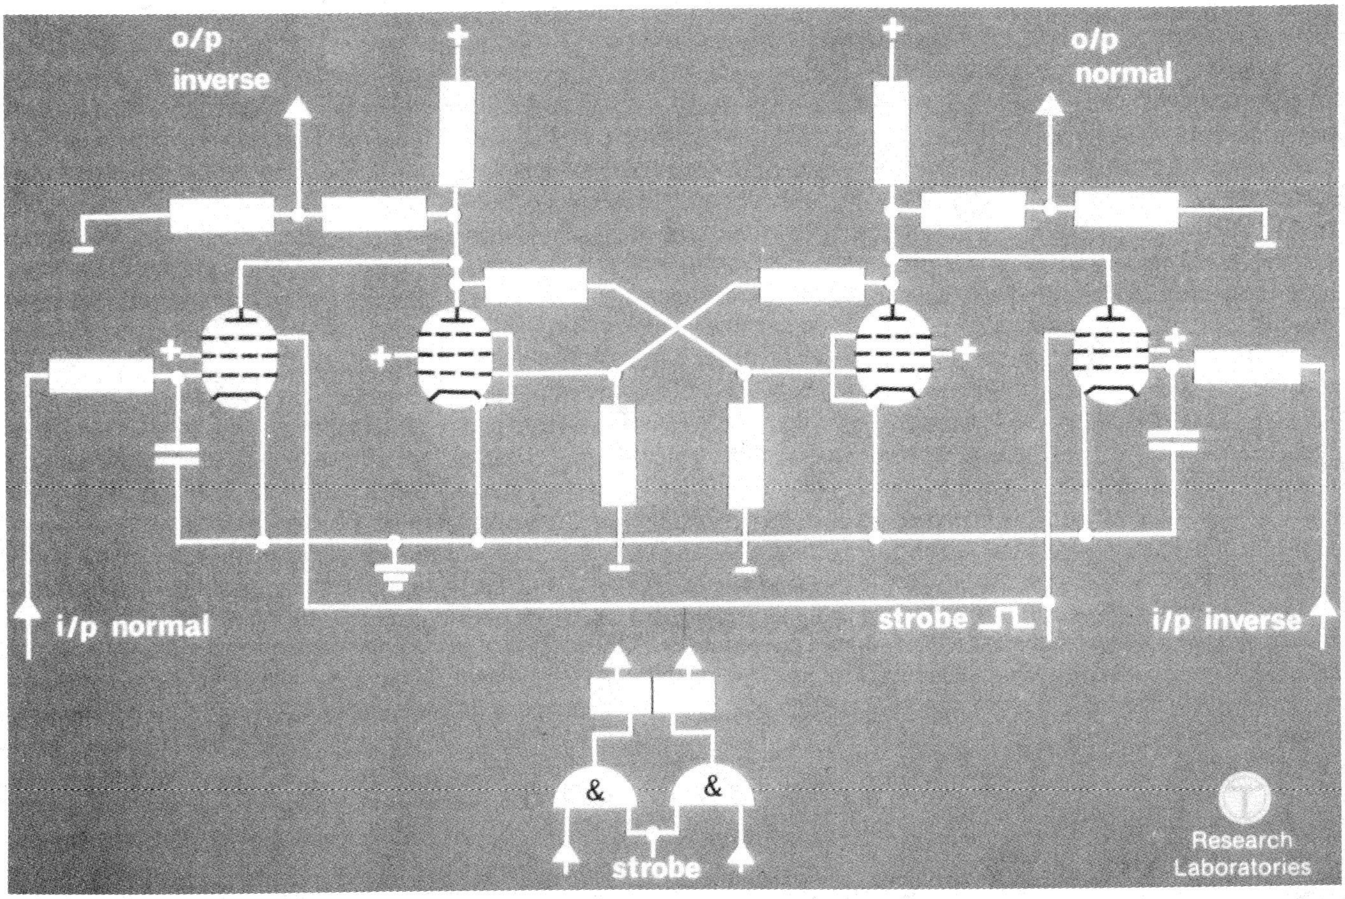

Figure 10. Eccles-Jordan one-bit store.

between the start and stop holes of the message tape, some rings receiving all the pulses and others such sprocket-hole pulses as the logic allowed to be transmitted to them. The rings ceased to count after the end of the message and were prepared for the next scan by clear and start-address pulses, as previously described. Thus as data were read off the message tape, the appropriately generated bit streams were available in synchronism with them.

The start-address switch for a ring was an electromechanical rotary switch with an arm that connected in turn to the cathodes of all the thyratrons in the ring. The contact on which this switch stood at any one time marked the position from which stepping was to start for the next scan of the message. This point was indicated to the output printer by wiring on additional banks of the rotary switch (not shown in Figure 7). The point could be constant for all scans throughout one program, or the point could be progressive, the switch being stepped one position at a time during each scan of the tape so that all the positions of one or more rings would be used as starting points during the course of one program of tests.

Another option to suit parallel processing was that the start-address switches should take not one but five steps for every scan of the tape. For each program the programmer had to select and arrange for one of these modes of operation to be effective. The actual stepping of the switches took place during the scan.

The simplified schematic diagram (Figure 8) shows the five separate bit streams c1 to c5 of the received enciphered message emanating from the message-tape reader together with streams from the bit-stream generator. Selecting streams from the bit-stream generator and adding them modulo-2 to the received message data would produce data streams m1 to m5, which would in some way reflect the clear text of the message if the rings had used the correct data and starting points. The machine was not, however, permanently wired to add bit streams as in the diagram. Instead, all the bit streams were taken to switches in a processor where they could be associated with function units and decade counters as decided by the cryptanalyst-programmer operating the machine. The function units included not only modulo-2 adders requiring three tubes as already shown, but also Boolean adders needing only two tubes (Figure 9), inversion units with one tube, and Eccles-Jordan 1-bit stores using four tubes (Figure 10).

The shift registers for the message-tape outputs consisted of 1-bit stores connected in series through simple delay networks. The counters used vacuum tubes that could not be made to count reliably in scale-of-ten; hence each decade was made up of one scale-of-two and one scale-of-five counter. The outputs from the counters were thus given on seven wires for each decade that had to be decoded by relays to decimal equivalents for the output printer. Cryptanalysts operated Colossus by setting up a program and running it, using the result of that program to set up the next, and so on, until all the required data were obtained.

Operation

The full schematic diagram (Figure 11) shows the construction and operation of the machine in detail. The sprocket-hole pulses that were produced continuously by the message tape provided basic timing for all the operations of the machine. In the master control, the sprocket-hole pulses were used to generate clock pulses CL1, which coincided with sprocket-hole pulses occurring between the start and stop pulses received from the tape reader, and pulses CL2, which followed the CL1 pulses with a delay of about half a digit time. The bit streams c1 to c5 from the message tape were stored in 6-bit shift registers so that six consecutive characters of the messages were available for parallel processing in five individual processors, each of which was supplied with two different consecutive characters of the message; for simplicity the diagram shows only one processor. The message digits were clocked into the shift registers by CL1 pulses, as were the output digits from the thyratron ring stores into 1-bit stores, by which means the message-tape digits and the ring digits were presented to the processor in exact synchronism.

There were switches in the processor by which the data streams could be connected to function units and the function units interconnected with a final connection to a gate circuit, which allowed time for all the processing for one character from the message tape to be completed before the result was clocked by a CL2 pulse to the output counter. The result left the total accumulated on the counter unchanged or added one unit to it. Four decade switches enabled the programmer to specify a number that was a set total for the program being executed. If during one scan of the message the set total was reached or passed by the counter, that fact was communicated to a "read-total" gate and to the master control.

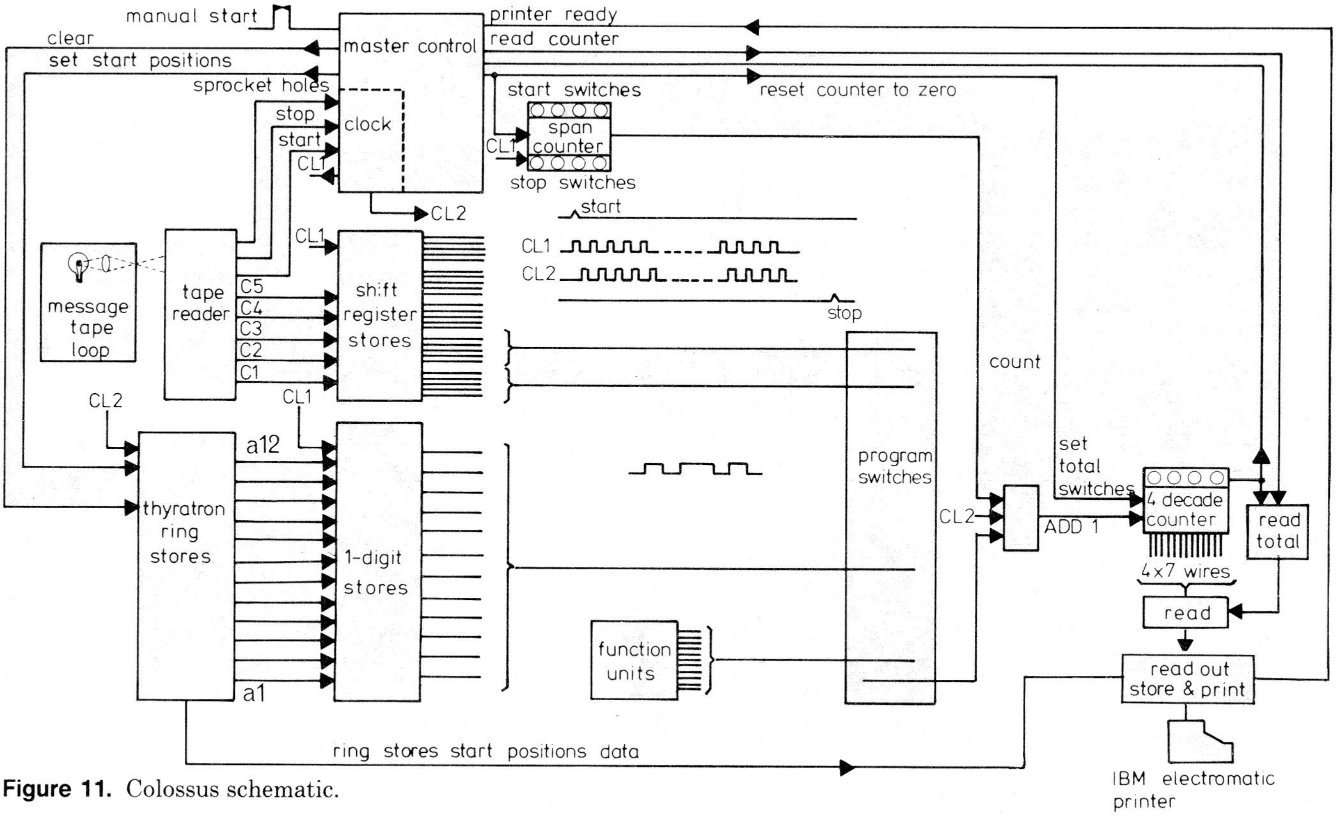

Figure 11. Colossus schematic.

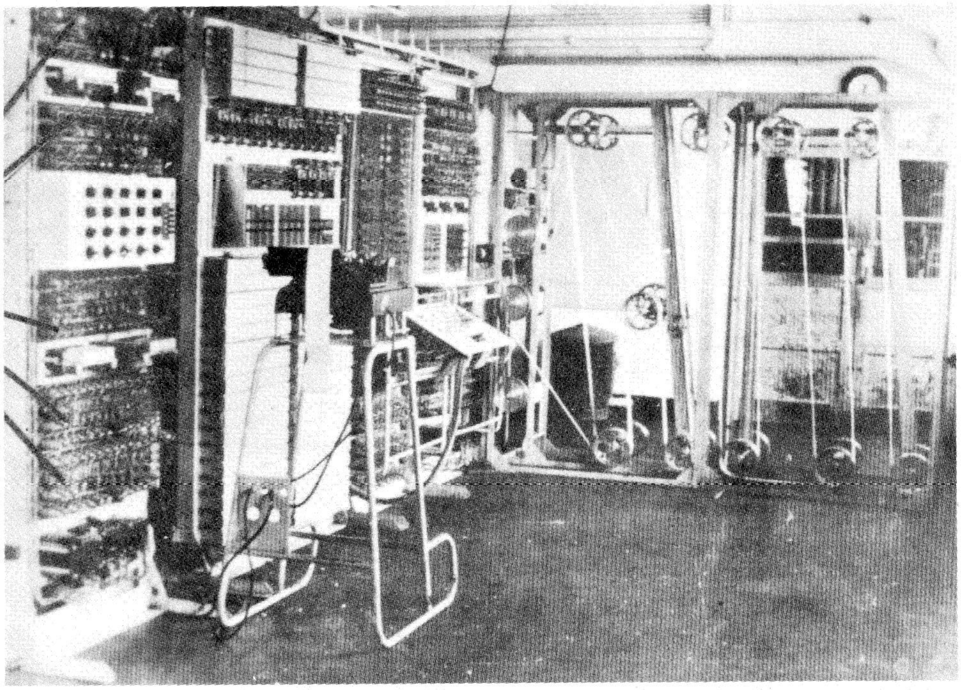

Figure 12. Front view of Colossus.

To run one program, the programmer selected the ring stores to be used and set them to operate in the required start-address modes. This operation also prepared the selected rings for the first scan of the message. The programmer then operated the program switches to set up the required program and the output counter switches to the set total he had chosen. More often the cryptanalyst-programmer in charge of the machine wrote the program in symbolic notation on a slip of paper and gave it to a Wren to execute. Each machine was operated by two Wrens who prepared and loaded tapes and attended to the machine, including setting up and running given programs.

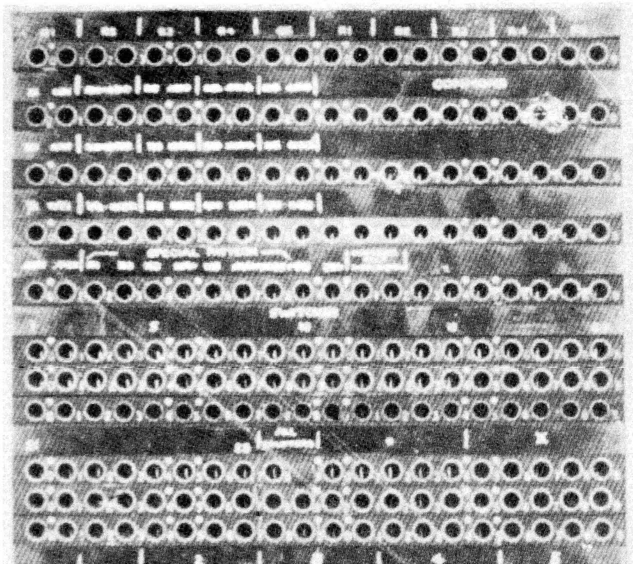

Figure 13. Colossus program switches.

When the program was set, the message-tape motor was switched on, and when the tape had reached running speed, a manual start switch was pressed. Processing started with the receipt by the master control of a start-hole pulse, and the CL1 and CL2 pulses following the start pulse caused the program to be executed. On the occurrence of the stop pulse—provided that the printer was ready and not still engaged in the last readout—the master control sent a "read-counter" pulse to a read-total gate. The output from that gate caused the output-counter total to be read out and stored on relays associated with the printer, together with the start positions of the ring stores at the beginning of the scan just completed. The read-total gate did not, however, produce output unless the set total was reached during the scan.

The set total was usually chosen to be a value that few of the scans would reach so as to minimize the quantity of data read and printed. If the set total was passed during a scan of the message but the printer was not ready to receive another readout of data when the stop pulse was received, the master control did not send the read-counter pulse and suspended all operations except the printout and tape revolution until it received a printer-ready signal and a stop signal from the tape. It then sent a test-counter pulse to the read-

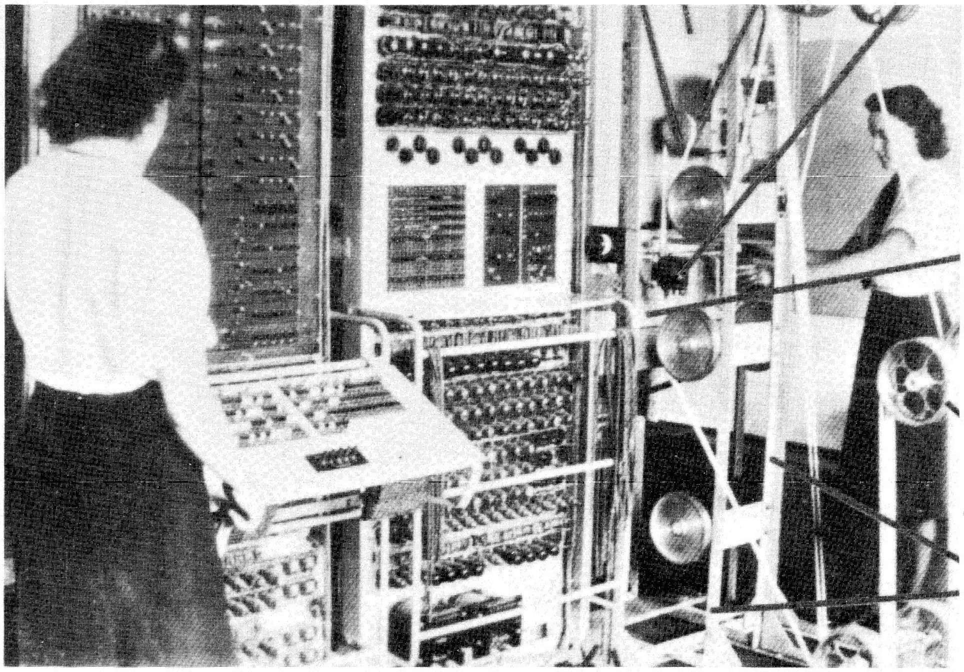

Figure 14. Controls of Colossus.

total gate that caused readout to take place and resumed normal operation: resetting to zero the output counter and a span counter (described in the next paragraph) and sending clear and start-address pulses to the ring stores to prepare them for the next scan of the message.

The master control program had to allow time for all of these operations to be completed without mutual interference, and the blank section of the message tape between stop and start holes had to be of minimum length of about 150 characters so that the master control operations could be completed before the start hole restarted the processing.

The processing was not always successful. If no result could be obtained from a message, corruption of the text in reception was a possible cause. Sometimes we tried processing not the whole message but just sections of it, using a span counter that counted all the CL1 pulses but allowed the output counter to operate only between two number outputs of the span counter set by switches on that counter.

As more machines became available, it became convenient to hardwire some of the much-used or more complicated programs instead of using the switches to set them up. For that reason there were individual variations between machines, although all were basically the same.

In Figure 12 a bedstead with its pulleys and two tapes can be seen. The two tapes were necessary for the Robinson machines; although Colossus needed only one tape, the two-tape design was continued so that while one tape was running, another could be loaded into the machine to save machine time between messages. Either tape could have a length up to 25,000 characters—just over 200 feet. Also visible is the automatic typewriter with printout on a narrow roll of paper and the decade switches for the set totals of five counters required for parallel processing.

The first Colossus used telephone plug and cord connections between jacks as switches to set up the programs (Figure 13). This feature was included in the production machines, but most of the programming was done by lever-type keys on a large panel especially

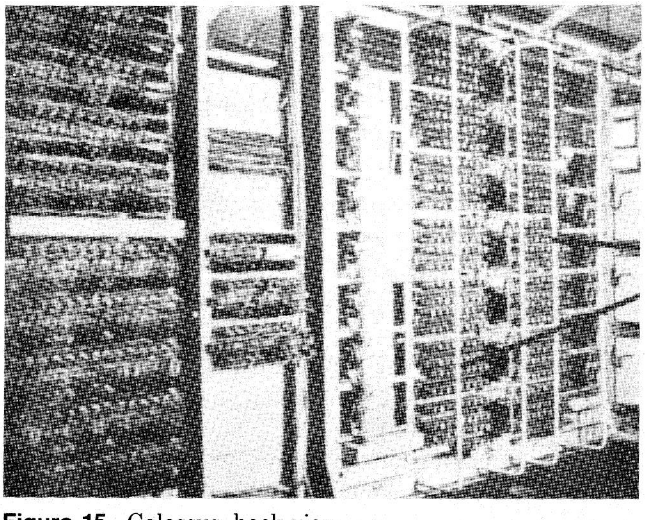

Figure 15. Colossus, back view.

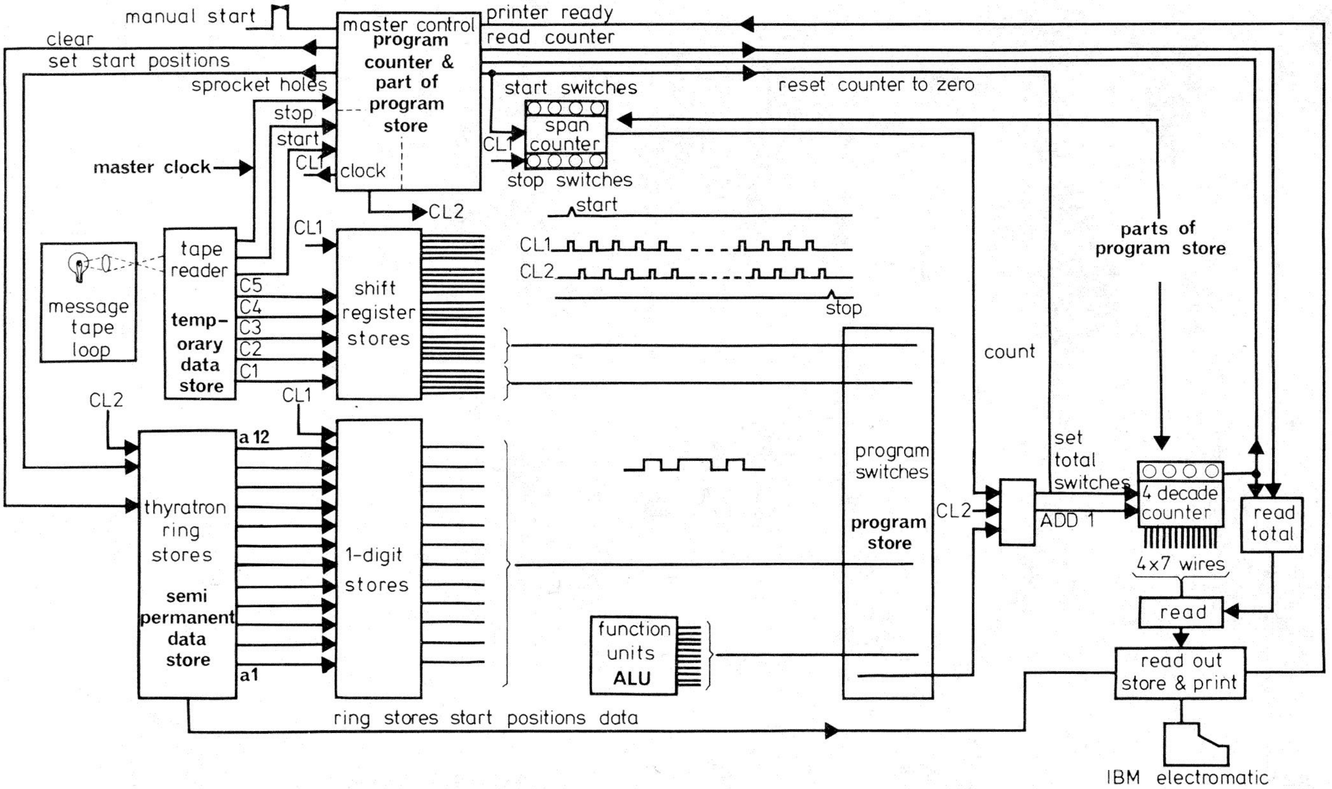

Figure 16. Colossus schematic with units labeled with computer equivalents.

designed for parallel processing. These can be seen in the photograph of the controls (Figure 14), which also shows decade switches for controlling the span counter. The back view (Figure 15) is not that of the line of racks pictured in the front view (Figure 12) but of a second line of racks parallel with the first.

Colossus had features now associated with digital computers—semipermanent and temporary data storage, arithmetic and logic units including branching logic, and variable programming—that may justify its being regarded as the first digital computer (see Figure 16). The control of Colossus by wired logic and of a computer by stored-program logic to make it a general-purpose machine is the most obvious difference. Colossus was a special-purpose machine that was required to attain the highest possible processing speed; the technology of the day would have compelled the use of wired logic even if stored-program logic had been thought of by then. Each switching operation took about 10 microseconds to complete, and the total program time between CL1 and CL2 pulses was much less than 200 microseconds. Nevertheless there was enough flexibility in the programming for several tasks for which the machine had not been specifically designed to be executed. Colossus was comparable in conception and in processing power with the ENIAC, which was designed as a computer but without stored-program logic and in operation some two years after Colossus.

Colossus undoubtedly made a contribution to the development of computers in Britain by showing Turing, Newman, and others what electronics could do, and that knowledge turned their minds to computers immediately after the war. I was not myself interested so much in computers as in telephone exchanges, and in 1946 I discovered that what was needed besides equivalence to make electronic exchanges economic was time-division multiplexing, but no practical use was made of the discovery until after I had retired—that is another story, however.

Conclusion

The events recounted happened so long ago that when I came to write this paper I found I had forgotten many of the details. The happy result was that the inquiries I had to make caused me to contact many of the people who took part at the time, some of whom I had not seen since. That was in itself a pleasure, and the information they gave me was a great help which I gratefully acknowledge. I am also indebted to the Post Office laboratories, no longer at Dollis Hill but on the east coast at Martlesham, for providing the illustrations used in the paper.