NCC PUBLICATIONS

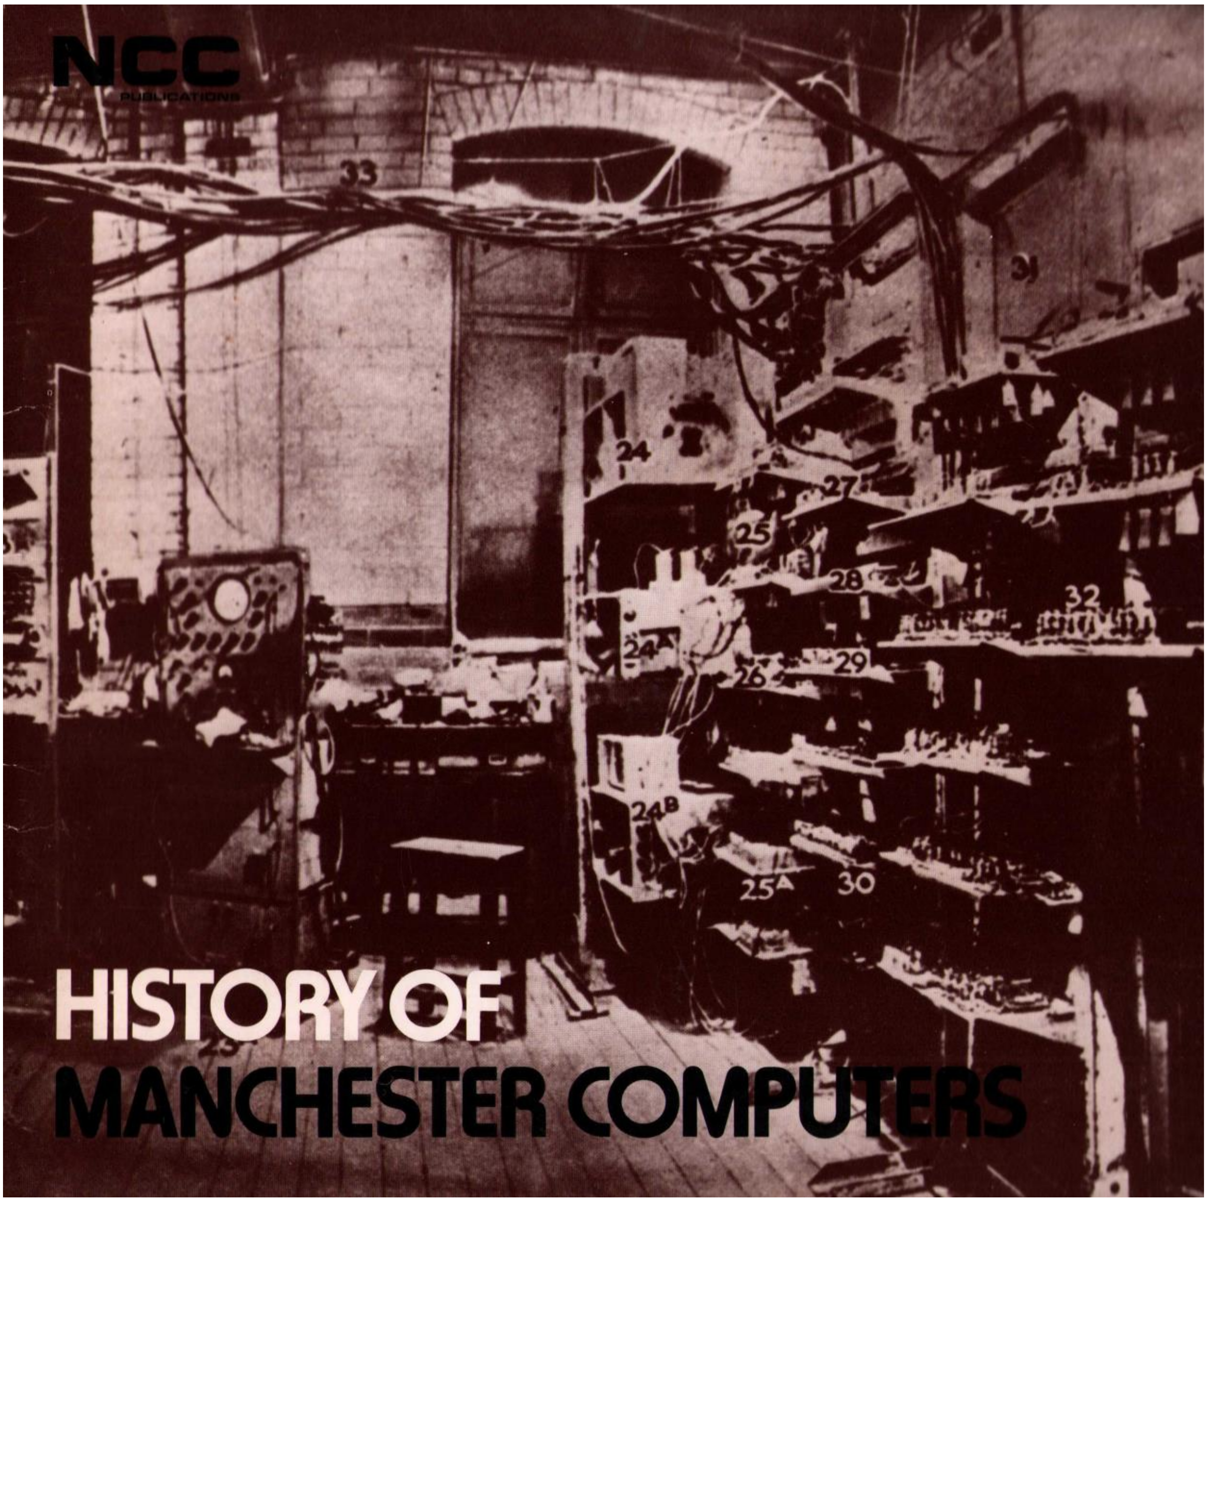

HISTORY OF MANCHESTER COMPUTERS

Contents

Keywords for information retrieval (drawn from the *NCC Thesaurus of Computing Terms*): Computer history.

First edition published 1975

Published and printed in the United Kingdom of Great Britain and Northern Ireland by NCC Publications, the National Computing Centre, Oxford Road, Manchester M1 7ED, England.
Distributed in the United States of America and in Canada by The Hayden Book Company Inc., 50 Essex Street, Rochelle Park, New Jersey

ISBN 0 85012 155 8

Set in 9 on 10 point Times using the NCC CAPRICORN computer-aided publishing system.

A history of Manchester computers

by
S H Lavington, MSc PhD MIEE MBCS

Preface

Computer design at Manchester University has been a continuous activity since 1946. This has resulted in the development of five prototype computers and, through collaboration with local industry, five commercially available derivatives. This is surely a unique record.

Although all five computers have constituted landmarks in the development of hardware and software, the first and fourth designs, known respectively as the Manchester Mark 1 and Atlas, are of special historical interest. These machines are therefore highlighted in the descriptions which follow.

The co-operation of Ferranti Limited and International Computers Limited is gratefully acknowledged in the production of the computers themselves and in this account of them. The author would also like to thank the following individuals for help in recalling historical details:

Prof T Kilburn, Prof D B G Edwards, Prof F H Sumner, and Prof D Morris, all of the Department of Computer Science, Manchester University
Prof D Aspinall, Prof R A Brooker, H J Crawley (National Research Development Corporation), T R Duffy, Prof R L Grimsdale, Prof M H A Newman, Miss J Pugh (Science Museum, London), Prof D Rees, Prof B Richards, Dr A A Robinson, N H Robinson (Librarian, Royal Society), Prof J S Rohl, J M F Scholes (National Research Development Corporation), C Somers (Ferranti Ltd, Archivist), B B Swann (formerly Ferranti Ltd), G Thomas, G C Tootill, prof M V Wilkes, Mrs C Williams (nee Popplewell), Prof F C Williams, Prof A M Uttley.

Simon Lavington
Manchester, August 1975

The Mark 1
Background

The word 'computer' prior to 1940 meant only one thing: a clerk equipped with a hand calculating machine who would 'compute' the standard calculations required for wages, actuary tables, gunnery, etc. From 1935-1945 the application of electronics to 'automatic computers' generally concentrated on developing faster calculators related to specific problems. Examples are the series of cryptanalysis machines built at Bletchley between 1940-45 for the British Foreign Office (*reference 1*), and the huge ENIAC machine commissioned late in 1945 for the US Army Ordnance Department for ballistics calculations (*reference 2*). By the end of the war there were thus groups in America, Britain and Germany working relatively independently on automatic electronic calculators. None of these early machines fitted Babbage's concept (*reference 3*) of a universal computer and it is clear that the designers were fully aware of the limitations of their projects. These limitations were partly concerned with input/output, partly with the nature of control-transfer (branching) facilities, but mainly with the lack of a suitable internal store.

The need for a storage device had also occurred towards the end of the war in connection with the processing of radar signals. By 1946 two lines of approach were being pursued: one based on acoustic waves propagated down a delay line (eg a column of mercury), and the other based on electrostatic charges stored on the phosphor screen of a cathode ray tube (CRT). In early 1946 the delay line technique seemed more practical, since CRT stores suffered from short retention periods due to charge leakage. In Britain, at the Ministry of Supply Telecommunications Research Establishment (TRE) – a centre renowned for electronic innovation – a group was working on CRT stores under Dr F C Williams. In the autumn of 1946 Williams made certain discoveries which resulted in the invention of a method of storing digital data on a CRT screen which differed fundamentally from previous attempts. The first digit was actually stored in November and, with an eye to computer applications already being discussed, the first patent application for 'Williams Storage' was filed by the Ministry of Supply in Williams' name on 11th December 1946. Within a few days of this F C Williams left TRE to accept the Chair of Electrical Engineering at Manchester University. He took with him a young member of the TRE group named Tom Kilburn.

Professor Williams was – and is – a highly inventive electrical engineer with a wide range of interests. He has made original contributions in many fields, apart from electronic pulse techniques and digital storage – for example, variable-speed ac motors, infinitely-variable mechanical drives, the velodyne. Although enthusiastic about building a small computer to verify the usefulness of his storage invention, Professor Williams' active involvement in computers only lasted from about 1946 to 1952, after which he devoted himself increasingly to other areas of electrical engineering. Tom Kilburn, on the other hand, developed a lasting interest in computer design. Graduating in Mathematics from Cambridge in 1942, he was immediately drafted to TRE where he learnt about electronics whilst working on radar and related projects. By 1946 he had attained a valuable combination of mathematical logic and practical electronics which was essential to the implementation of early computers. (Indeed, at least one contemporary group's progress (NPL) was inhibited by difficulties in reconciling the views of the 'mathematicians' and the 'engineers'.) The move to Manchester by Williams was natural – he was a graduate of the University and had been on the staff before the war. The move for Kilburn was unexpected: he was anticipating returning to his old Cambridge college, but Williams arranged for him to be seconded by TRE to Manchester to complete work for a PhD. This secondment lasted for two years and had the valuable side-effect of allowing components from TRE to be supplied to the Manchester computer project. It should be remembered that all materials were scarce in those days and

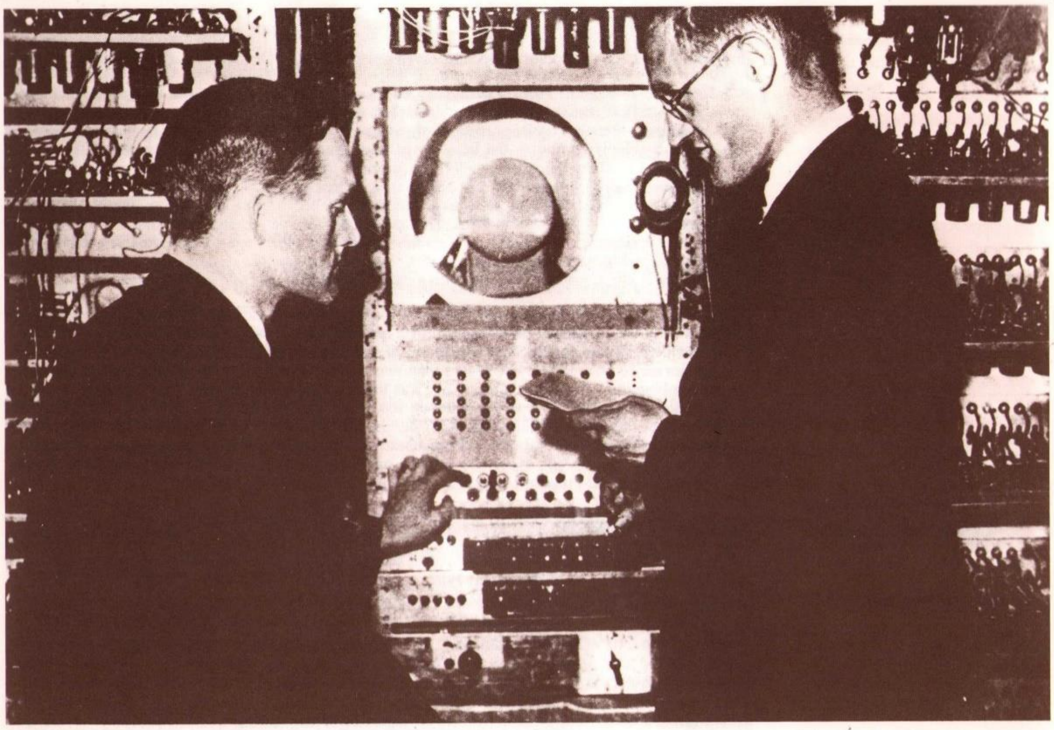

Tom Kilburn (left) and F C Williams shown at the control panel of the Manchester University Mark 1 computer. The prototype first ran a program on June 21st, 1948 and is believed to be the world's first stored-program computer.

even the wooden TRE packing-cases were eagerly shared out amongst the technicians as domestic firewood during the 1947 fuel crisis!

Although the Williams/Kilburn computer project did not depend on an appointment to any particular University, its arrival in Manchester proved fortunate. Professor M H A Newman had taken up the Chair of Pure Mathematics at Manchester in October 1945, with I J Good and D Rees joining the Mathematics Department at the same time. All three had worked as mathematicians at the cryptanalysis centre at Bletchley during the war. Max Newman had at one stage led the team which produced, in December 1943, the first of the COLOSSI series of electronic 'computers' (*reference 1*). Professor Newman had his own understanding of the theory of universal computers (*reference 4*) and in February 1946 applied to the Royal Society for a grant "for a projected calculating machine laboratory at Manchester University". The application was enthusiastically endorsed by a Royal Society investigating committee whose members included Professor D R Hartree (Cambridge) and Professor P M S Blackett (Manchester). The grant, approved in July 1946, was at the rate of "£3,000 a year for five years for salaries, together with the sum of £20,000 to be spent on construction during the same period" (*reference 5*). Professor Newman later explained the context of this Royal Society application in a report dated October 1948 (*reference 6*):

> There are two distinct fields of inquiry connected with large automatic computing machines. These are first, the engineering problem of designing the machine, and especially of finding a satisfactory 'memory' unit, or method of storage of information; secondly, the mathematical and logical problems of finding the best use of such machines and investigating their effect on the development of mathematics itself. The original intention was that the laboratory should be concerned mainly with problems of the second, or mathematical type, and that a machine should therefore be constructed as soon as possible, modelled on one or other of those which were already under construction in 1946. The most suitable type seemed to be that which was being made at Princeton under John von Neumann's direction. The memory unit, [the Selectron] which it was hoped to have ready in six months, was very compact in design, and the idea was to buy from the Radio Corporation of America a sufficient quantity to build into our machine.
>
> It was at this time that Professor Williams was appointed to the Chair of Electrical Engineering. He was already engaged on the design of a memory unit, based on different principles from any previously contemplated. Since his work had only recently begun it was not reasonable to expect his storage to be ready for use in less than two years. Nevertheless, such a fortunate opportunity to combine the two sides of the enquiry was too good to be missed, and it was decided to rely on Professor Williams' storage in spite of the delay that appeared to be involved. What has happened, however, is that while the American unit is still not quite ready, owing to unforeseen difficulties, a small prototype of the machine using Professor Williams' storage came into action about three months ago in the Electrical Engineering Laboratory, and is thus the first of these automatic general-purpose computing machines to have actually worked.

In 1946 Newman asked Rees to select some problems in group theory, etc which might be suitable for solution by computer. He also sent him to America to attend the Moore School seminars in July/August 1946 (see later). Rees wrote an informal report of these seminars which he passed to A M Turing at NPL upon his return. Rees has said that they did not at that time realise the difficulty of programming a universal computer and that his own and Good's interest soon faded. Good left Manchester in 1948 and Rees in 1949. The Royal Society grant of £35,000 remained substantially intact for several years, eventually providing for the construction of a building to house the University's Ferranti Mark 1 computer in 1951.

Although Williams and Newman did not meet until the autumn of 1946, it is possible that the interest of Newman and Professor Blackett (Physics) in computers was a factor in the appointment of Williams to the Chair of Electrical Engineering. Newman and his mathematicians took no active role in the design of Manchester computers, but he did provide a general enthusiasm for the project including arranging the acquisition of war-surplus Post Office racks from Bletchley. He gave the specification for some of the earliest 'useful' programs on number theory, but his personal interest in using computers seems to have disappeared by about 1950.

Before describing the first Manchester computer in more detail it is appropriate to review briefly the wider climate in which Williams and Kilburn started work. Generally speaking most people in 1946 regarded the idea of a 'universal' computer as something of a scientific oddity, with little practical significance. In America, the ENIAC group was striving to change this opinion. Thus in the summer of 1946 a course was held at the Moore School of Electrical Engineering at the University of Pennsylvania, in which outline plans were put forward by Eckert and Mauchly and others (*eg reference 7*) for the design of a successor to ENIAC. The new machine, generally referred to as EDVAC, was to be a universal stored-program computer. F C Williams, whilst in the USA on a radar assignment in 1946, had seen the ENIAC at the suggestion of Professor D R Hartree who was advising the Moore School at the time. Hartree, of Cambridge, had a long-standing interest in automatic calculating machines (*references 8, 9*). Williams has said that his main impression of ENIAC was that it demonstrated how an assembly of more than 18,000 thermionic tubes could be made to stay error-free for long enough to get useful work done. Most people then were dubious about the reliability of electronic logic and some groups, for example TRE (*see below, and reference 10*), went to great lengths to incorporate self-checking facilities throughout their computer. M V Wilkes, who attended the Moore School seminars, produced in May 1949 the Cambridge University EDSAC (Electronic Delay Storage Automatic Computer) inspired by the EDVAC ideas. EDSAC was about a year ahead of the various American EDVAC-derived projects.

Although the Moore School reports constitute a concise, readable account of the problems and possible solutions they do not represent the only source of ideas from that period. The concept of a universal stored-program computer was by 1946 being talked about in a number of centres in America, Britain and Germany and the modern computer would probably have been developed by some means or other from any one of these centres, if the others had not existed.

None of the Manchester group and few people in the United Kingdom generally were in direct contact with ENIAC or the EDVAC designers. An outline of some of the British computer projects in progress during the period, apart from the Williams/Kilburn effort (*see references 4, 11, 12*), is as follows:

National Physical Laboratories (NPL)

Initial investigations begun in 1945. A M Turing, who had published his early ideas on universal computers in 1936 (*reference 13*), joined NPL from Bletchley in October 1945. By 1947 the ACE design group was under way using delay-line storage. The ACE Pilot Model finally worked in the autumn of 1950.

Manchester (Newman)

Plans for a computing machine laboratory put forward in March 1946 (as described previously). In due course this project became merged with the Williams/Kilburn machine.

TRE

Initial planning begun early in 1946. After Williams and Kilburn left for Manchester A M Uttley designed TREAC, a parallel, self-checking computer using Williams' Tube storage and a drum backing store. TREAC was completed by early 1953 and remained in use for several years.

Cambridge University

EDSAC design under way by November 1946: completed May 1949. M V Wilkes' team included W Renwick as one of the principal engineers.

Birckbeck College, London

Dr A D Booth started work in 1947 and in the same year completed most of the arithmetic section of his Automatic Relay Computer (ARC) using Siemens' high speed relays. It was built at the laboratory of the British Rubber Producers Research Association, Welwyn Garden City. A magnetic drum store for ARC was still under development in 1949. Several other small electronic computers followed.

These British groups worked for the most part independently of each other, as may be judged from the considerable differences in technology, order code format, and central processor logic of the various machines. There were also differences in emphasis. Thus, at Cambridge the aim was to construct a machine which, while not being of the highest technological performance possible, would be good enough to enable them to do programming development. Subsequent to 1949 this development concentrated on the provision of a software library of mathematical subroutines. Wilkes took the view (*reference 14*) that "when a machine was finished, and a number of subroutines were in use, the order code could not be altered without causing a good deal of trouble. There would be almost as much capital sunk in the library of sub-routines as the machine itself, and builders of new machines in the future might wish to make use of the same order code as an existing machine in order that the sub-routines could be taken over without modification". At Manchester the emphasis was on the continual enhancement of technology and processor architecture. Dr A A Robinson, a research student at the time, has said that the atmosphere at Manchester in 1947 was like that of a nineteenth-century inventor's workshop with regard to the surroundings and the enthusiasm. The Mark 1 was built, significantly, in the Electrotechnical Laboratories and, in fact, three versions of the machine were commissioned between 1948 and 1949. Folklore has it at Manchester that 'F C' never wrote a program in his life and that Tom wrote just one – the world's first!

The enhanced Mark 1 in June 1949. The original photograph appeared in the illustrated London News under the following heading: 'A marvel of our time: the 'memory' machine which can solve the most complex mathematical problems'. The key to the numbers in the picture was given as follows:

Power supply distribution rack
1. Power supply meters
2. Main power switches

Basic tuning rack
3. Master oscillator (controls timing of all operations on machine)
4. and 4a. Horizontal deflection (for horizontal scan of CRTs)
5. Test gear

Vertical deflection rack
6. Cathode-ray tubes selection circuits (for selecting CRT reading and writing)
7. CRT vertical deflection circuits

Operator's controlling position
8. Controlling circuits
9. Prepulse generator (initiating process of obeying each instruction)
10. Erasing circuits (these automatically erase the previous information as new information is received in the 'Memory')
11. Monitor CRT (this makes visible the contents of any of the six storage CRTs)
12. Loading keys (used in the manual insertion of numbers in the 'Memory')
13. Selection switches (these select the position where the numbers are loaded into the machine)
14. Operating controls
15. High voltage power supplies

Pulse separator rack
16. Controlling circuits
17. Pulse separator (this counts the pulses received from the master oscillator)

Controlling circuit rack
18. Controlling circuits
19. 'Stop' circuit (automatically stopping the machine at the end of the calculations)
20. Controlling circuits
21. Erasing circuits (have the same function as defined in 10)

No 2 Control circuit rack
22. Controlling circuits
23. Horizontal deflection for the monitor CRT (horizontal scan)
23a. Fault-locating portable test gear

Storage rack
24, 24a, and 24b. Cathode ray tube storage circuits (this is really the 'Memory' and consists of Williams' Tubes in metal boxes)

Regeneration rack
25, and 25a. Regeneration circuits automatically and continuously 'refreshing' the machine's 'Memory'
26. Subtractor (subtracting incoming numbers from the numbers already stored in the 'Memory')

Computer
27. Transfer gates (regulating the passage of numbers from place to place in the machine)
28. Monitoring circuit (this circuit transmits the information stored in any of the storage CRTs and displays it on the monitor CRT)
29. Selecting circuit (this selects the computing circuit to be used)
30. Adder (adds an incoming number to the number already stored in the 'Memory')

Multiplier racks
31. Storage CRT (used during multiplication)
32. Multiplying units
33. Interconnecting wires

The prototype machine

Upon arrival at Manchester in December 1946, Williams and Kilburn set about perfecting a digital store, at first using the commercially available type CV1131 12-inch diameter cathode ray tubes (*references 15, 16*). The principle of a two-state electrostatic store can be visualised from the following simple experiment. Start with a focussed CRT beam and turn the beam current on (thus producing a charged 'dot') and off again repeatedly. *Negative* voltage pulses will be induced by capacitive coupling in a pick-up plate placed close to the outer surface of the CRT screen. Now move the beam whilst it is on so as to write a 'dash' on the screen, then move the beam back whilst the current is off, and then switch on the current again. This time a *positive* voltage pulse is induced. With dots and dashes representing logical 0's and 1's, readable as negative and positive voltage signals, a binary storage system is available. Other representations such as a 'focus/defocus' system were also used. Now although the electrostatic charge leaks away in about 0.2 seconds, automatic refreshing (re-writing) of the information in less than 0.2 seconds is a simple matter electronically. (cf a modern MOS solid-state store.) Since the refresh rate is rapid, long term drifts in electrode supply voltage, etc are not critical and a robust store can be made from standard components. In contrast, the mercury acoustic delay-line stores chosen by other workers had to be constructed to close physical tolerences. The biggest advantage of the CRT store was that it allowed random access whereas other contemporary systems were sequential.

By the Autumn of 1947 the Manchester group had successfully stored 2048 digits for a period of hours (*reference 15*) and the way was clear for the construction of a prototype computer 'to subject the system to the most searching tests possible' (*reference 17*). Kilburn took the initiative with the logical design. The 'baby machine', as it was called, had a specification which may be expressed in modern terminology as follows:

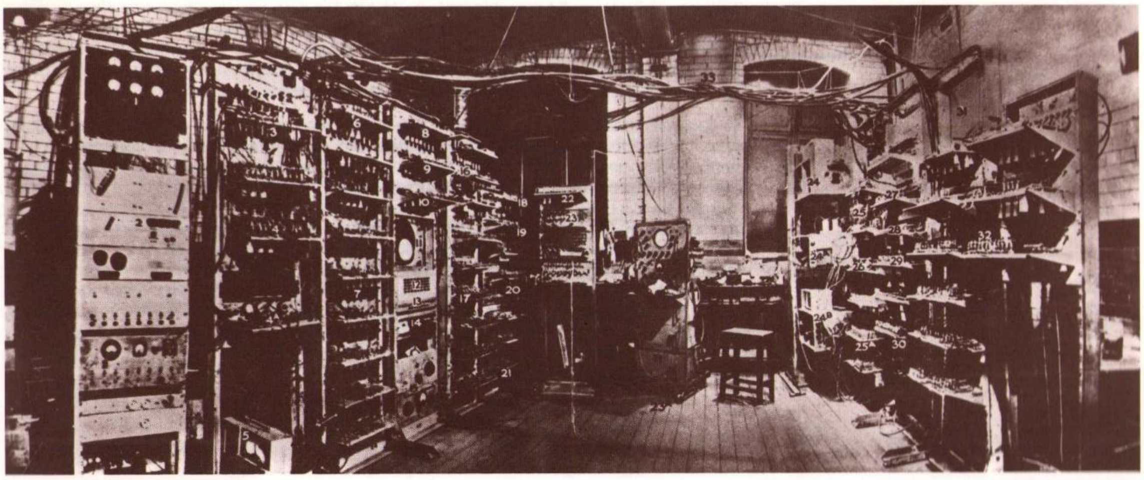

32-bit word length
serial binary arithmetic using two's complement integers
single-address format order code
main store: 32 words, extendable to 8192 words, random access
computing speed: 1.2 milliseconds per instruction.

The instruction format had three bits assigned to the function field, 13 bits to the address field and the remaining 16 bits were unused. The main store consisted of a single CV1131 Williams Tube, with each 32-digit line occupying about 10 cms on the screen and being scanned in 272 microseconds. A complete 'beat' of 306 microseconds consisted of 32 × 8.5 microsecond digit periods plus a four digit fly-back time. The rhythm of the whole processor was synchronised to this store beat. There was notional provision for extending up to 256 Williams Tubes to yield a total storage capacity of 8192 words. The arithmetic unit was based on a serial subtractor and the logic employed EF50 pentode tubes, used widely for wartime applications. Using this technology, flip-flops (bistable circuits) were extremely costly and temporary storage throughout the central machine was implemented with Williams Tubes wherever possible. Thus the accumulator and control register (instruction counter) were Williams Tubes. One incidental advantage of the use of CRT's was that the contents of main store, accumulator and control register could be viewed on a monitor CRT during or after a computation – so providing a simple output mechanism. Input for the prototype was via a 32-position keyboard and operator's control switches.

The machine first ran a program in June 1948 (*reference 18*) and as far as can be ascertained it was therefore the world's first stored-program computer. A complete diagram of the prototype Mark 1 is given in reference 17 and figure 1 is a simplified version showing the main flow of information. The Williams Tube which implemented the control register was also used to hold the present instruction (PI) itself subsequent to its being read out of main store. Either the value of control or the value of this PI could be fed from the 'control' Williams Tube to an 8-bit (extendable to 16-bit) flip-flop register known as the staticisor. This staticised function bits (F) and operand addresses (S) during the execution of an order, and then staticised the address of the next order during the instruction-fetch phase. An interesting anomaly was that numbers were stored with the least significant digit on the left – a system which makes sense to engineers if not to mathematicians!

As has been noted, two's complement was used to represent negative numbers, though the now familiar rules for addition/subtraction and the formation of the complement of a number were not implemented in the usual way. The main emphasis of the project at this time was to prove the practicability of the Williams Tube for realising the stored-program concept and so the arithmetic logic was kept as simple as possible. The subtractor was the *only* arithmetic facility provided, it being preferred to an adder because a subtractor can be used without alteration to form complements and to perform additions whereas the converse is not true. As may be seen from figure 1 an operand entered the accumulator by being fed through the subtractor; this 'subtract from zero' thus complemented the operand before it reached the accumulator. The effect of a positive load could be programmed in two orders by performing a negative load and then subtracting the contents of the accumulator from a zero operand. An addition of two quantities p, q was programmed by a four-order sequence as follows:

load negative p
subtract q
store the resulting $(-p - q)$
load negative to achieve $(+p + q)$

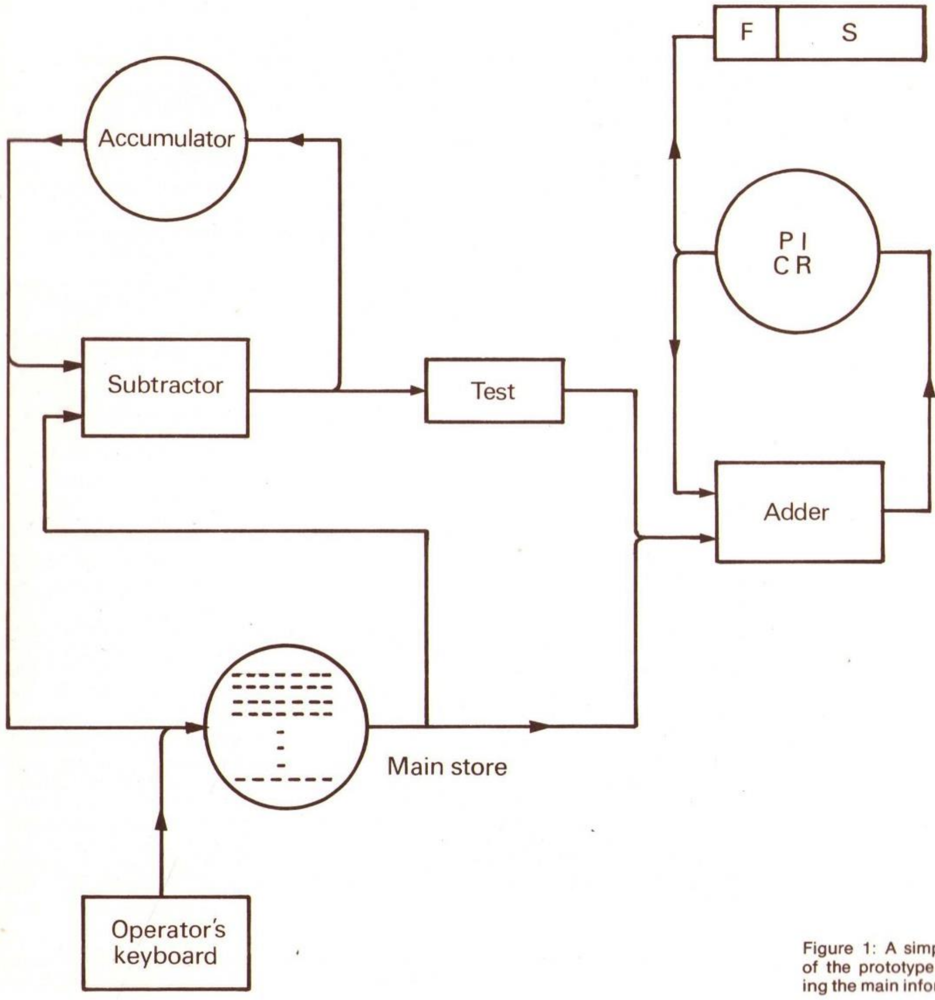

Figure 1: A simplified diagram of the prototype Mark 1 showing the main information paths

The order code of the prototype had provision for eight functions, specified as in table 1.

decimal value of function bits	an early notation	modern mnemonic	explanation of operation
0	s, C	JMP S	Absolute indirect unconditional jump: set the control register equal to the contents of address S*
1	c+s, C	JRP S	Relative indirect unconditional jump: add the contents of address S to the control register*
2	−s, A	LDN S	Load negative: set the accumulator equal to the negated contents of address S
3	a, S	STO S	Store: copy the contents of the accumulator to address S
4 or 5	a−s, A	SUB S	Subtract: set the new value of the accumulator equal to the former contents minus the contents of address S
6	Test	CMP	Compare against zero: the value in the accumulator is tested. If it is less than zero, one is added to the control register thus causing the next sequential instruction to be skipped
7	Stop	STP	Stop: cease automatic mode, and await manual commands from the operator's keyboard

* Note that +1 was always added to the control register at the end of every order, so the programmer used JMP and JRP to point to an instruction one before the line he intended to jump to

** To economise on logic elements only partial decoding of the function bits was carried out

Table 1: The order code of the prototype Mark 1.

Three demonstration programs were run on the prototype machine, the first one involving determination of the highest factor of an integer by a method which would give a long run, the result of which could be easily checked. To quote reference 18, "the highest proper factor of 2^{18} was found by trying in a single routine every integer from $2^{18}-1$ downward, the necessary divisions being done not by long division, but by the primitive process of repeated subtraction of the divisor. Thus about 130,000 numbers were tested, involving some 3.5 million operations. The correct answer was obtained in a 52-minute run. The instruction table in the machine contained 17 entries."

The original program was written by Tom Kilburn. G C Tootill, an engineer on loan to Manchester from TRE from mid-1947 to mid-1949, also wrote programs for the prototype and a notebook kept by him over the period 4th June to 28th November 1948 has survived. From entries in this notebook it seems that Kilburn's program was first run on Monday 21st June. Williams has since described their relief at seeing the correct answer! (*reference 19*) The original program has been lost but Tootill's notebook gives an amended version written by him on 18th July, 1948. This is reproduced in figure 2, in which Tootill used C to represent the 'computer' (ie accumulator) and Cl to represent the control register. The other letters a, b, and r were respectively the number for factorising, a trial factor, and a remainder.

Another program, run in mid-July, involved a long division routine written by A M Turing (and corrected by Tootill). Turing was appointed to the nominal post of 'deputy director of

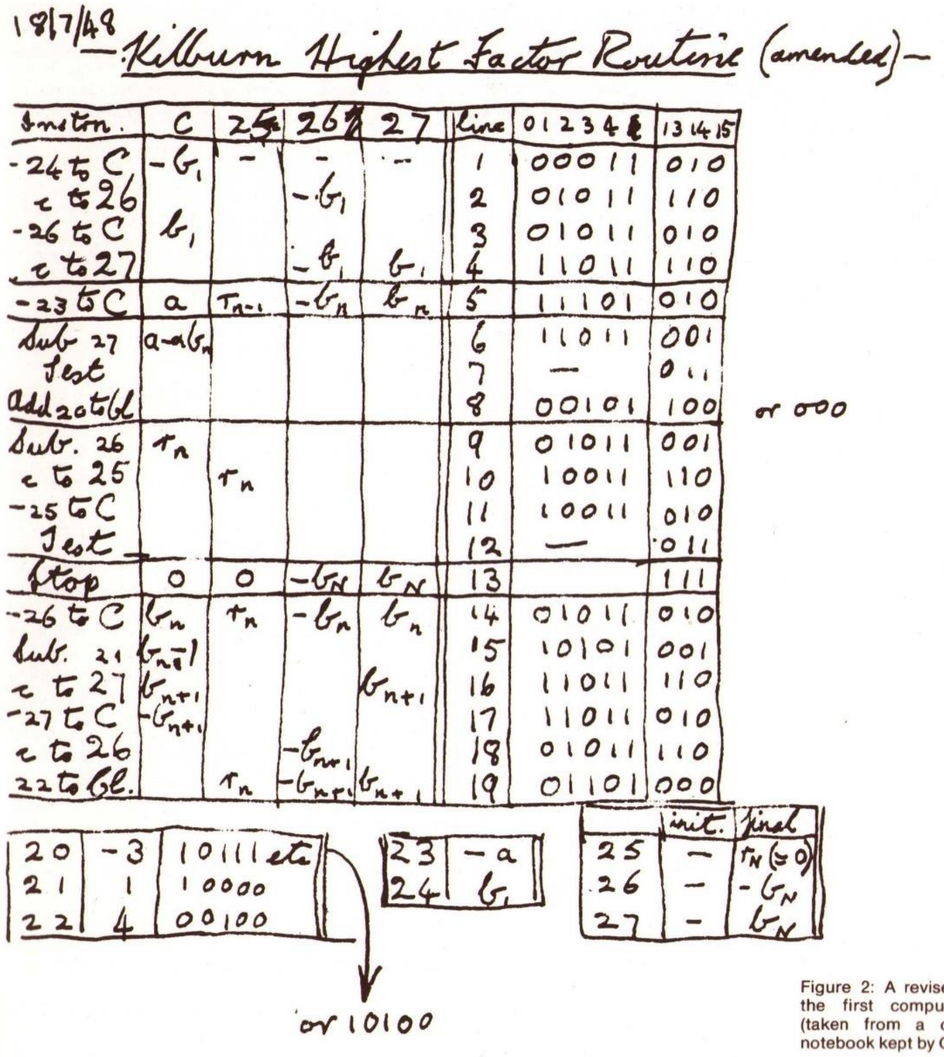

18/7/48 — Kilburn Highest Factor Routine (amended) —

Instrn.	C	25	26	27	Line	0 1 2 3 4	13 14 15
-24 to C	$-b_1$	–	–	–	1	00011	010
c to 26			$-b_1$		2	01011	110
-26 to C	b_1				3	01011	010
c to 27			$-b_1$	b_1	4	11011	110
-23 to C	a	r_{n-1}	$-b_n$	b_n	5	11101	010
Sub 27	$a-ab_n$				6	11011	001
Test					7	—	011
Add 20 to Cl					8	00101	100 or 000
Sub. 26	r_n				9	01011	001
c to 25		r_n			10	10011	110
-25 to C					11	10011	010
Test					12	—	011
Stop	0	0	$-b_N$	b_N	13		111
-26 to C	b_n	r_n	$-b_n$	b_n	14	01011	010
Sub. 21	$b_n - 1$				15	10101	001
c to 27	b_{n+1}			b_{n+1}	16	11011	110
-27 to C	$-b_{n+1}$				17	11011	010
c to 26			$-b_{n+1}$		18	01011	110
22 to Cl.		r_n	$-b_{n+1}$	b_{n+1}	19	01101	000

20	-3	10111 etc
21	1	10000
22	4	00100 (or 10100)

23	−a
24	b_1

	init.	final
25	–	r_N (= 0)
26	–	$-b_N$
27	–	b_N

Figure 2: A revised version of the first computer program (taken from a contemporary notebook kept by G C Tootill)

the Computing Machine Laboratory', with the status of a Reader in the Mathematics Department under Professor Newman, on 29th September, 1948. (His salary was the first call upon the Royal Society grant.) Turing had resigned from NPL during the summer (*reference 20*) and before arriving in Manchester had written asking for details of the prototype computer's order code. He did not actually become involved in using the Mark 1 until the spring of 1949 and the 'Computing Machine Laboratory' was not a physical reality until 1951.

The enhanced Mark 1

From about August 1948 onwards the prototype was under intensive engineering development in order to provide a more realistic computing facility (*reference 35*). The design team was increased by taking on two very keen research students – D B G Edwards and G E Thomas. Two re-designed versions of the computer were commissioned in 1949. The first, working by about April, had the following specification (*references 21, 22*):

40-bit word length; two 20-bit instructions per word
serial arithmetic, with a double-length accumulator facility
two modifier registers (B-lines)
single-address format order code with a 26-function repertoire including hardware multiply
128 words of random access main store
1024 word drum backing store, 30 milliseconds revolution time
instruction time: 1.8 milliseconds – except for multiplication which was operand dependent.

The main store consisted of two improved 6-inch Williams Tubes giving 64 words per CRT. The drum, on which Thomas worked, was synchronised to the main central processing unit clock (*references 23, 24*), thus allowing extension to multiple drums. Fixed block transfers of 64 words took 30 milliseconds plus the latency for the block to arrive under the read/write head, the transfers being initiated via a manual switch. The drum, or "magnetic wheel" as it was first called because of its squat shape, contained novel features including phase modulation recording – still sometimes referred to as 'Manchester notation'. Perhaps the most significant feature of the Mark 1 was its inclusion of *B-lines* (modifier, or index, registers) – a Manchester invention. The two B-lines were stored on a Williams Tube, the symbol B being used simply because A (accumulator) and C (control) had already been assigned. One of Edwards' tasks was the implementation of the B store. Although still somewhat of an engineering prototype this version of the Mark 1 performed useful work on Mersenne Primes during the summer of 1949 (*reference 25*), at one point giving an error-free run of nine hours (June 16/17th, 1949). Newman specified the algorithm for the Mersenne calculations, and Kilburn and Tootill wrote the program. Turing later wrote a much faster program which he called the 'Mersenne Express'.

By October 1949 the Mark 1 had been enhanced again in the following respects (*reference 26*):

- extended order-code including a peripheral control order for programmed drum transfers and input/output transfers
- provision of input/output routines using 5-bit teleprinter code and paper tape devices.

This machine did useful work in 1949/50 including investigation of the Riemann hypothesis and calculations in optics (*reference 25*). It was closed down in the summer of 1950, by

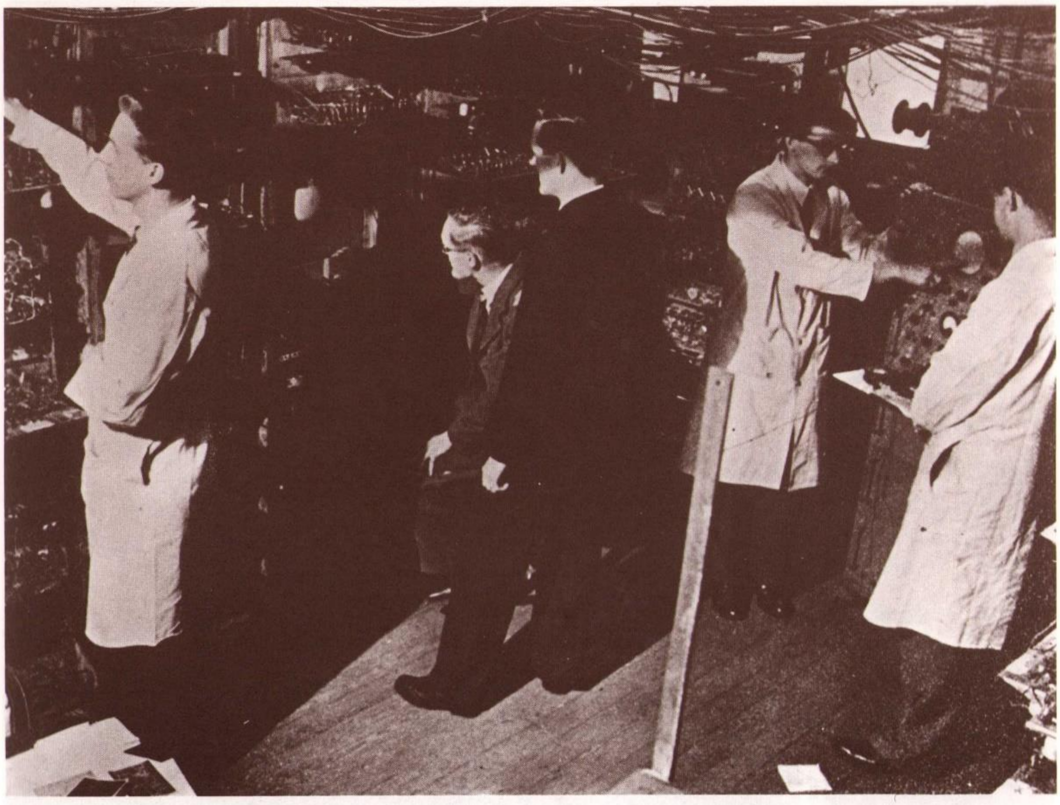

Some of the Mark 1 design team: (left to right) D B G Edwards, F C Williams, T Kilburn, A A Robinson, G E Thomas

which time the first production Mark 1 manufactured by Ferranti was nearing completion (see later).

During the three-year period late 1946 to late 1949 the active design team had averaged about four people, supported by the equivalent of about two technicians, including the fabulous Ida Fitzgerald on loan from TRE. Thirty-four Mark 1 patents were taken out via the National Research Development Corporation (NRDC) during 1946 to 1949. (The NRDC was actually in business from July 1949 onwards and patents previous to this were taken out by the Ministry of Supply and subsequently transferred to NRDC.) Since many of the ideas were joint efforts, two or three inventors' names usually appear on each patent application. A rough idea of the relative design contribution of the people involved may be gained from the analysis of the 34 patents given in table 2. Professor Newman's name

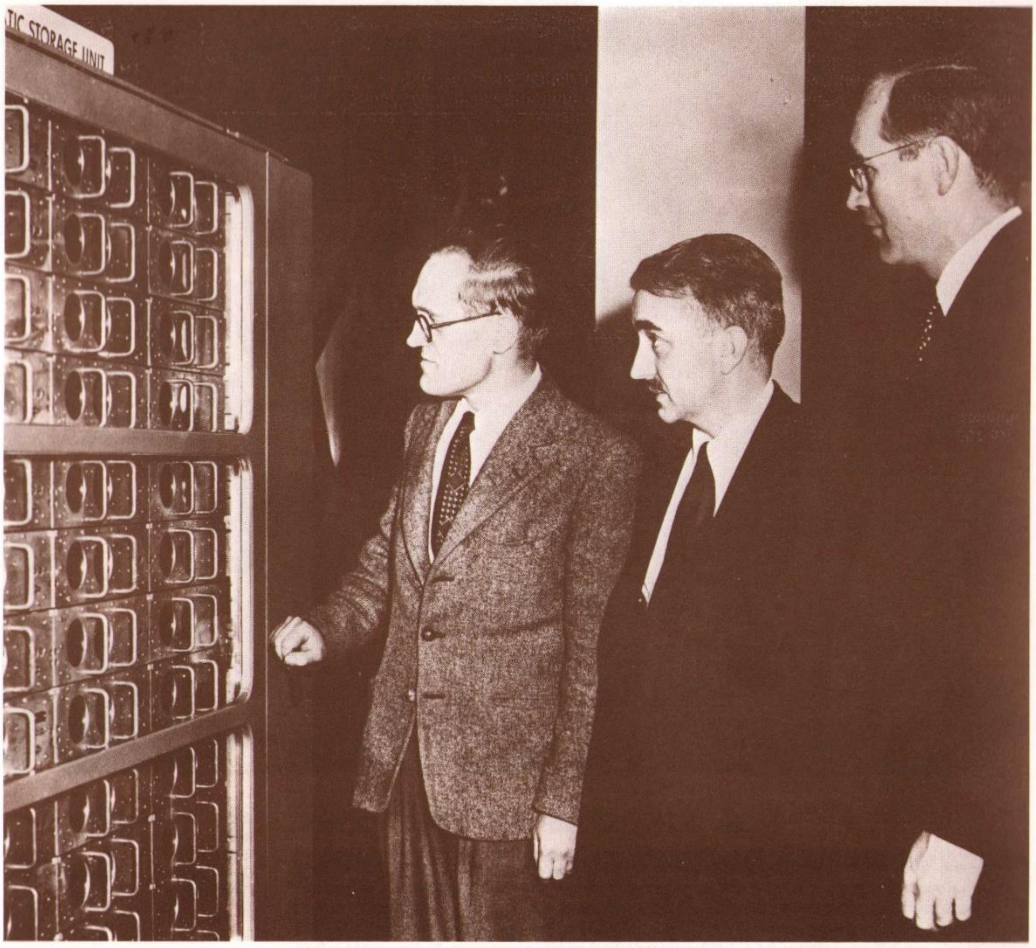

(left to right) F C Williams, H J Crawley (NRDC) and McPherson (IBM) seen looking at Williams Tube electrostatic storage units as used under licence by IBM in their 701 and 702 computers from about 1953 onwards.

name	dates at Manchester	number of patents of which an inventor or co-inventor
F C Williams	Dec '46 onwards	32
T Kilburn	Dec '46 onwards	27
G C Tootill	July '47 – April '49	12
A A Robinson	April '47 – April '49	7
G E Thomas	Oct '48 – Sept '55	3
D B G Edwards	Oct '48 onwards	2
M H A Newman	Oct '45 – Sept '64	1
J C West	Oct '46 – Dec '57	1

Table 2: An analysis of the inventors responsible for 34 Mark 1 patents, filed between 1946-49 (with acknowledgement to NRDC)

features in the famous B-modification patent along with Williams, Kilburn and Tootill. The contribution of J C West concerned the drum servo system.

Compared with the only other stored-program computer working in 1949 – the Cambridge EDSAC – the Manchester Mark 1 was much less of a finished product from the user's viewpoint but offered a comparable basic number-crunching performance (the Mark 1 in fact had a rather larger storage capacity). It excited interest amongst the general public, as may be judged from the series of photographs (see pages 7 to 15) which appeared in the Illustrated London News for June 25th, 1949. It was popularly christened MADM – Manchester Automatic Digital Machine – though this name was never used in the University.

Ferranti and NRDC involvement

The Manchester computer project quickly came to the attention of officialdom. As early as July 1948 Sir Henry Tizard saw the machine and "considered it of national importance that the development should go on as speedily as possible, so as to maintain the lead which this country has thus acquired in the field of big computing machines, in spite of the large amount of effort and material that have been put into similar projects in America" (*reference 6*).

The wider application of the Mark 1 design stems from October 1948 when Sir Ben Lockspeiser, at that time Government Chief Scientist, made an unannounced call at the Electrical Engineering Department whilst visiting Manchester University to see Professor Blackett. Lockspeiser was so impressed with the prototype computer he saw working there that he immediately initiated a government contract with the Manchester firm of Ferranti Limited to make a production version of the machine. The original letter from Lockspeiser to Eric Grundy, manager of the Ferranti Instrument Department, is reproduced in reference 19. It resulted in a five-year contract running from November 1948, involving an estimated £35,000 per annum (*reference 27*). This produced the commercially-available Ferranti Mark 1 and Mark 1 Star series of computers, the first of which was installed in the University in 1951 (as described later). More importantly perhaps, the contract established a fruitful link between the University and the computer industry which has been maintained to the present day.

Two Ferranti engineers seen at the console of the first production Ferranti Mark 1, installed at Manchester University in February 1951. This is believed to be the world's first commercially available computer.

Concurrent with this Ferranti involvement the Manchester Mark 1 was also becoming known in America. IBM were sufficiently interested to invite Professor Williams on an all-expenses paid two-week visit to the USA in July 1949. Williams recalls that, having accepted the invitation, he was already on board the Queen Mary just about to sail when an urgent telephone call arrived from Lord Halsbury of NRDC informing him that NRDC would pay all his expenses and that he was not to accept a penny from IBM! The National Research Development Corporation was formally set up on 25th June, 1949, after much concern that British ideas such as the jet engine were being exploited abroad to the eventual disadvantage of British industry. Professor Williams was NRDC's first customer (*reference 27*). Williams did not, of course, succumb to IBM and the visit went ahead with due propriety.

There is one immortal anecdote of this trip that should be recounted. Williams was shown round the IBM World Headquarters in New York and was fascinated to see the exhortation THINK emblazoned on doors and desks throughout the building. Williams then gave a talk

about the Mark 1 to assembled IBM staff. During question time one of them asked him how it was that a small team of two men and a dog had been able to produce a machine at Manchester which was ahead of anything produced by the mighty resources of IBM. "Well", said Williams, "we had an idea for storage and then pressed on regardless and built a computer round that store without stopping to *think* too much". An amazed gasp greeted this reply, followed by shrieks of hysterical laughter when the audience realised that their boss was not in the room.

Another view of the Ferranti Mark 1 at Manchester University. The cabinets contain logic doors such as the one shown in the next photograph.

Twenty years of technology: a Mark 1 logic door containing eight flip-flops as compared to an Atlas printed circuit board containing six flip-flops (upper hand) and an MU5 module containing four flip-flops (lower hand). The circuits are representative of about 1951, 1961, and 1971 respectively and the speeds of operation are approximately 200 nanoseconds, 20 nanoseconds, and 2 nanoseconds.

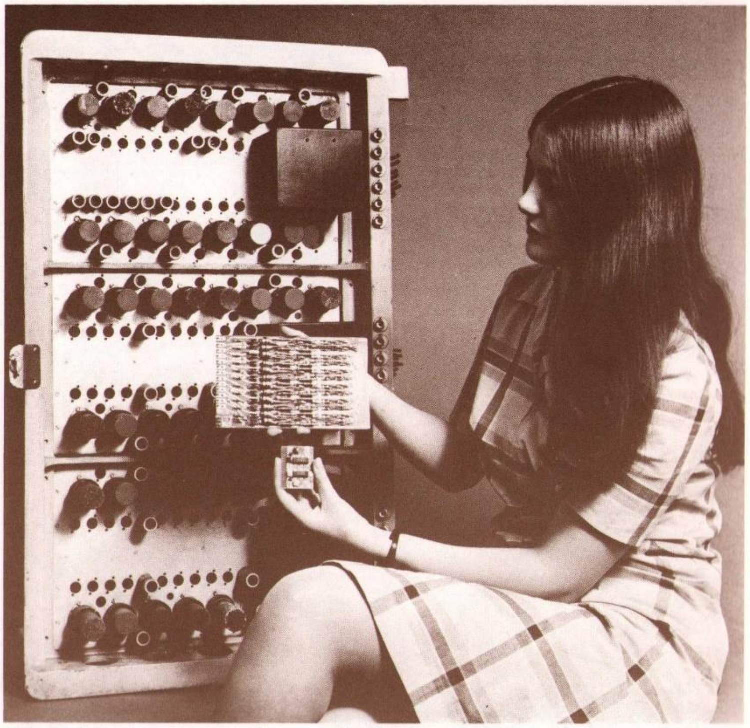

IBM later used the Williams Tube and other Manchester patents under licence (the IBM 701 and 702 computers had Williams Tube storage). A mutually beneficial agreement was arranged by NRDC – option agreement dated July 1949; licence agreement dated April 1951 (*reference 27*).

It is worth mentioning at this point the strenuous efforts made by NRDC between 1949 and 1951 to encourage an alert and competitive British computer industry. These efforts were not altogether successful and in the light of hindsight it appears that the industry was too cautious in the early days. From Manchester's viewpoint, involvement with NRDC chiefly concerned the administration of patent applications, extension of coverage overseas and the undertaking of legal action against infringement. For example, in the early 1950's a complex interference action was set up involving NRDC (Williams), Remington Rand (Eckert-Mauchly), Raytheon (Hergenrother) and the US Government (Haeff). By the end of 1956 NRDC held a total of 201 computer patents, including 81 originating from Manchester University, 43 from the National Physical Laboratory, 36 from various other

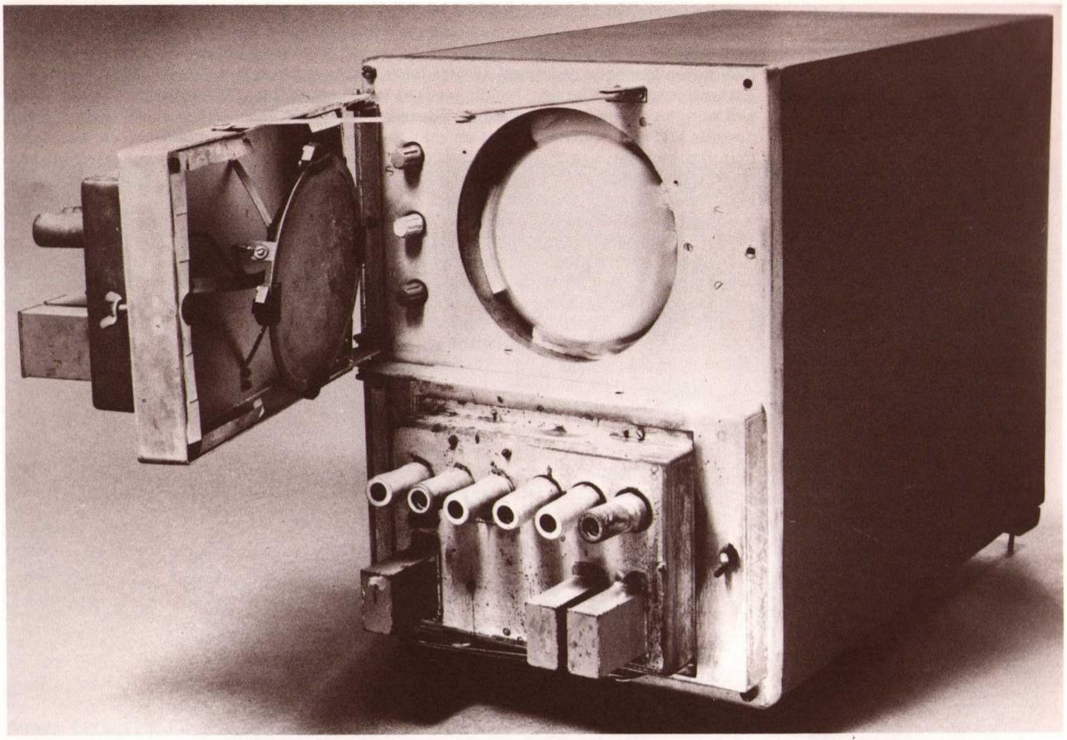

A twin Williams Tube electrostatic storage unit from a Ferranti Mark 1 Star computer. Each tube stored 32 × 40-bit words. The sense amplifier and pick-up plate have been swung clear of the upper CRT to reveal the screen of the tube. The Williams Tube was the first random-access digital storage device.

government establishments under the Ministry of Supply, 16 from Ferranti, 10 from NRDC itself, and the remainder from seven smaller sources. A revenue sharing agreement was set up, under which Manchester University had received £27,000 by the end of 1956. The burden of drawing up patent specifications and general liaison was ably undertaken by H J Crawley of NRDC, via frequent visits to Manchester. NRDC also gave assistance to the University in providing a number of computer scholarships and in underwriting certain annual expenses of the Computing Machine Laboratory in the early 1950's. (In the event, annual earnings payable to the Computing Machine Laboratory by outside users always eliminated the need to call upon NRDC finances to meet running costs.)

To return to the Ferranti contract, the University and Ferranti collaborated closely on the development of the production Mark 1. G C Tootill spent four months with Ferranti from August 1949 and by the end of November the logic design was substantially finished (*reference 28*). A A Robinson joined Ferranti in April 1950 and the computer was commissioned on the factory floor by the end of 1950. Meanwhile a special computer building was being planned at the University since the existing physical embodyment of the Royal Society Computing Machine Laboratory – a 20 feet square room, whose brown tiled decoration has been described as 'late lavatorial' – was totally inadequate for housing the Ferranti machine. The new building was to be financed from the 1946 Royal Society grant and architects were approached in December 1948. By July 1949 the necessary building permits had been obtained – not an easy task in post-war austerity Britain. It was hoped that the building would be complete by August 1950, but staff did not move in until mid-January 1951.

The Ferranti Mark 1 was delivered in February 1951 and although its rèliability during the first few months was somewhat questionable (*reference 29*) it was giving reasonable service by the Inaugural Conference (July 9th–12th, 1951). Ferranti sold a second Mark 1 (delivered 1952) to Toronto University where it was sometimes known as the FERUT computer. Under a contract placed by NRDC seven modified machines known as Mark 1 Stars were delivered between 1953-1957 (*reference 30*), two of which were exported (to Holland and Italy). As far as can be established the Ferranti Mark 1 (February 1951) was the world's first commercially-available computer, marginally ahead of the first UNIVAC machine (June 1951) which was descended from the Eckert-Mauchly BINAC. The LEO (Lyons Electronic Office) computer – the commercial manifestation of the Cambridge EDSAC – was working by December 1951.

The Ferranti Mark 1, referred to in sales literature as the "Manchester Electronic Computer" and by the maintenance engineers as the "blue pig", was similar to the October 1949 University Mark 1 except for the following points:

extended function repertoire, including B-line arithmetic
eight B-lines (20 bits each)
main store: 256 × 40 bit words on eight Williams Tubes
drum backing store: 3.75K words, extendable to 15K
add-time: 1.2 milliseconds, multiply time: 2.16 milliseconds

The fast multiplier, designed by A A Robinson, used a parallel technique which took up nearly one quarter of the computer's 4,050 thermionic tubes (*reference 31*). Then, as now, people argued about how much emphasis should be placed upon multiplication in assessing computer performance. Measurements performed on an early sample of Mark 1 jobs estimated that 16% of computing time went on drum transfers, 28% on multiplication and 56% on other arithmetic operations (*reference 35*). A contemporary benchmarking exercise rated the Mark 1 at about the same raw power as the NPL ACE (*reference 31*), even though the latter had a digit period ten times as fast. The favourable performance of the Mark 1 was attributed to its random access main memory and its fast multiplier. By any scale of judgement it must have been amongst the most powerful computers in the world at that time. A minor curiosity of the order code was a hardware random number generator (*reference 33*). One lighthearted use of this feature was in the 'creation' of Mark 1 love-letters, of which the following is an example (*reference 34*):

Darling Sweetheart,
You are my avid fellow feeling. My affection curiously clings to your passionate wish. My liking yearns to your heart. You are my wistful sympathy: my tender liking.
Yours beautifully
M.U.C.

A stored pattern on a Mark 1 Williams' Tube, using the focus-/defocus method for representing binary 0's and 1's. Note how the 40-bit words are divided into 5-bit characters for programming convenience. The extra 20-bit line at the top left of the picture gives the address of this 'page' when held on the drum backing store. This was the germ of an idea which later led to page address registers and virtual-to-real address translation on the Atlas computer.

Software development

The Ferranti Mark 1 provided a computing power far in excess of the University's own needs in 1951 and so outside bodies also came to make use of the facility. This was officially encouraged by the Ministry of Supply, the nominal owners of the machine, on the grounds of helping to disseminate knowledge about the applications of computers as well as providing a design tool for industry. The first regular outside users appeared in the Spring of 1952 and were charged at the rate of £20 per hour. The service included the

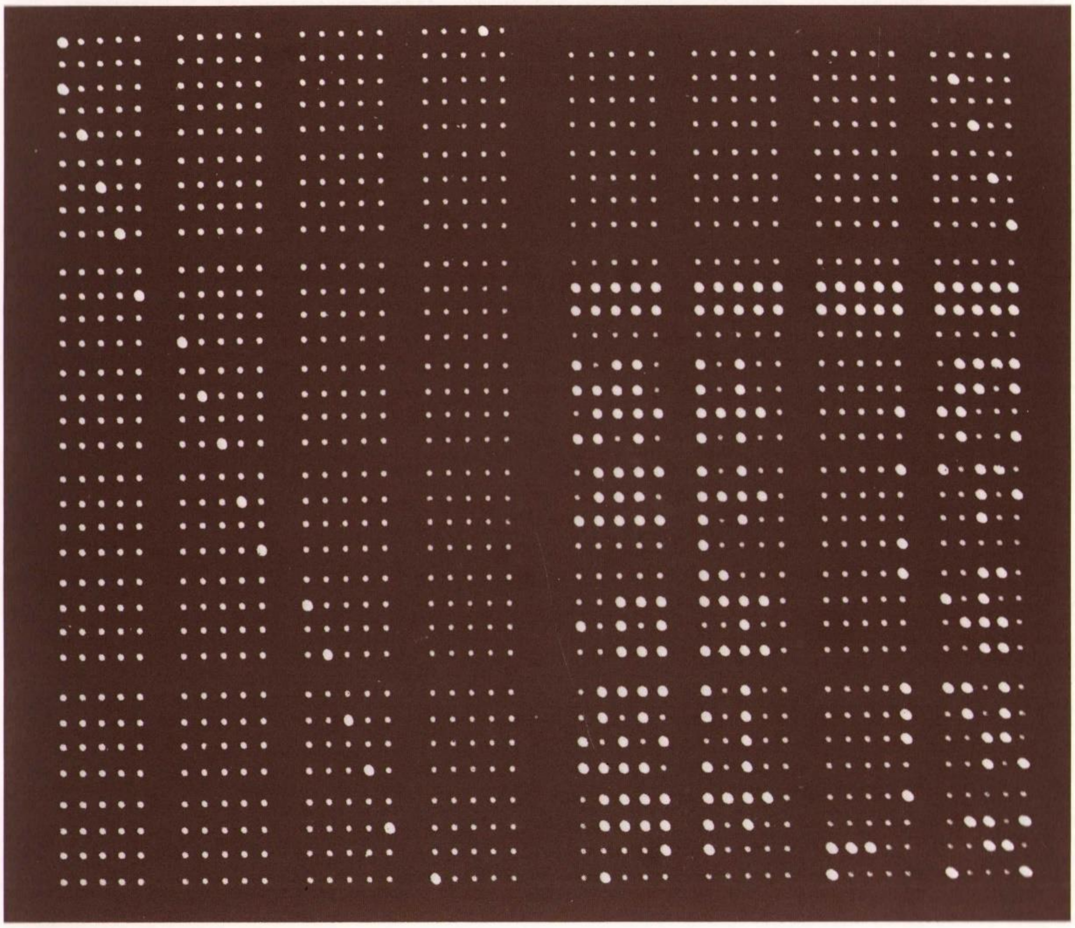

teaching of programming and help with debugging. An analysis for the year 1955 (*references 36, 37*) shows a user community made up from the following sources:

9 departments of Manchester University
6 departments of other Universities
3 research associations
9 government establishments
7 industrial firms doing engineering design

During the year 104 people were trained to use the machine and 66 scientific papers were published based on machine results. Of the 100 hours average useful computing per week in 1955 about 12 hours were allocated to Computing Laboratory staff, 30 hours to other University departments, 13 hours to other academic institutions and 45 hours to industry and government departments. The running costs of computer maintenance, staff salaries and equipment were partly met out of earnings and partly from contributions from general University funds and from the Department of Scientific and Industrial Research (DSIR) who had taken over nominal ownership of the Mark 1 from the Ministry of Supply in 1952. The Computing Machine Laboratory earnings amounted to about £20,000 in 1955. Some of this went on running costs and the remainder was used to support research in machine design and to build up a fund to permit eventual purchase of a better machine. In 1954 the Mark 1 was transferred to the spacious new Electrical Engineering Department in Dover Street, thus permitting the hardware and software research to proceed under the same roof.

All this activity required the support of system software and a user service, which can be said to have been established in a small way from October 1949 with the appointment of Cicely Popplewell to the Computing Machine Laboratory. It was anticipated that Turing would take an interest in software development and indeed he wrote the first Manchester programming manual (1950), based on what was later called 'Scheme A' input. He took little part in subsequent enhancements because by 1951 his energy was turning to morphogenesis – the 'growth and form of living things' (*reference 38*) – in connection with which he used the Mark 1 for the solution of partial differential equations. Turing died suddenly in June 1954. The Mark 1 software development gained momentum subsequent to the arrival in October 1951 of R A Brooker. His Mark 1 Autocode system (*references 39, 40*) was available by March 1954 and is worth describing in some detail since it appears to have pre-dated any FORTRAN implementation. (The first FORTRAN user manual, entitled "The FORTRAN Automatic Coding System for the IBM 704 EDPM" was dated October 1956.)

The purpose of the Autocode system was to provide a simple, easy-to-learn programming language for scientific users having small or medium-sized problems. In addition, one of the implementation objectives was to simulate a one-level store so that the user did not need to organise his own drum transfers. It was possible to achieve this one-level store on the Mark 1 in a reasonably balanced way because the access time for reading an operand from the drum (about 40 milliseconds) happened to be about the same time as a floating-point addition via an interpretive library routine. When running Autocode programs the first 128 tracks on the drum were reserved for instructions and the 40-digit locations on the remaining 128 tracks were reserved for variables labelled v1, v2,.... v8192. Individual routines were transferred to the fast CRT store as they were required. To gain access to a variable an interpretive routine in fast store first determined on which track it lay, then transferred that track or 'page' to the fast store, and finally selected the particular

TREAC, the parallel self-checking computer designed at the Telecommunications Research Establishment between 1946-53. At the bottom right are 12 Williams' Tubes, each storing two bits of a 24-bit word for a 1024 word parallel high-speed store. TREAC had a relatively simple order code. The photograph shows the computer as it was in 1955.

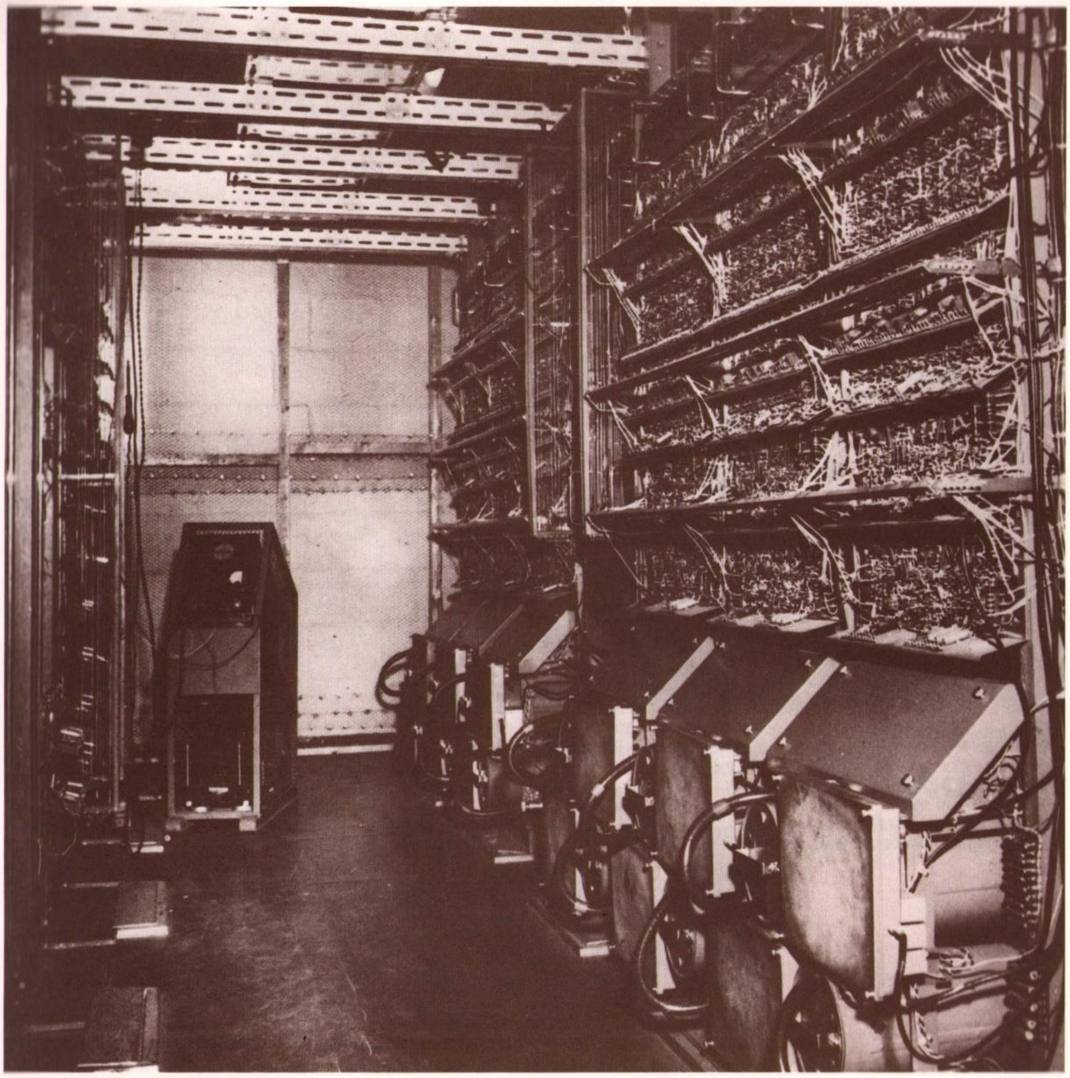

'line within the page'. Since successive operands were quite often located on the same page, unecessary drum transfers could be eliminated (*reference 40*).

In the Autocode system, arithmetic was normally performed on floating-point variables v1, v2,... etc with provision for integers n1, n2,... to be used as indices and counters. Simple conventions also existed for control transfers, intrinsic functions, input/output and simple job control using symbols from the five-track teleprinter code. An impression of the neatness of the system may be gained from the following Mark 1 Autocode sequence which prints the root mean square (RMS) of the variables v1, v2,... v100. (note: the symbol * causes printing of a variable to ten decimal places on a new line, and F1 signifies the intrinsic function 'square root'):

```
n1 = 1
v101 = 0
2v102 = vn1 × vn1
v101 = v101 + v102
n1 = n1 + 1
j2, 100 ≥ n1
v101 = v101/100.0
*v101 = F1(v101)
```

Negotiations for the replacement of the Mark 1 were begun with Ferranti in November 1955, the replacement machine eventually starting useful work in February 1958 (see below). Maintenance on the Mark 1 was discontinued in December 1958. After lengthy negotiations with the DSIR and the Science Museum at South Kensington, the Mark 1 was dismantled in June 1959 and given to Ferranti to form the basis of a projected company museum. In the event, the remains of the machine were disposed of by Ferranti some years later.

Several well-known names were associated with the Computing Machine Laboratory software over the seven-year life of the Mark 1, amongst which should be mentioned F H Sumner who joined the group in September 1957 having used the machine a few years earlier as a chemistry research student. Sumner became part of the design team of subsequent Manchester computers.

Meg and Mercury

Returning to the hardware development, Kilburn's group began work on a Mark 2 computer, often known as the megacycle machine Meg, in 1951. Meg ran its first program in May 1954. The design aim was to produce a faster, more compact, easier-to-maintain version of the Mark 1 which for user convenience would include floating point arithmetic facilities. This was achieved (*reference 41*) as follows. The digit period of the serial arithmetic unit was shortened to one microsecond and *parallel* access was made to the store in 10-bit short words at a time. Floating point hardware was provided using a 30-bit mantissa, a 10-bit exponent and a base of two. The floating point add time was 180 microseconds and floating point multiplication took 360 microseconds using a parallel paired-multiplicand technique. Eight 10-bit B-registers were provided, B-arithmetic orders being executed in 60 microseconds. Instructions were 20 bits long with the same single-address format as the Mark 1. Since magnetic core stores had not become generally available in 1952 the Meg prototype used a CRT main store, with provision for switching to cores for the production version. The use of semiconductor ('crystal') diodes helped to reduce the number of thermionic cathodes by 57 per cent as compared to Mark 1. This and the use of the smaller type 6CH6 pentode tubes also produced a reduction in power consumed (12kW instead of Mark 1's 25kW and in the physical size of the processor.

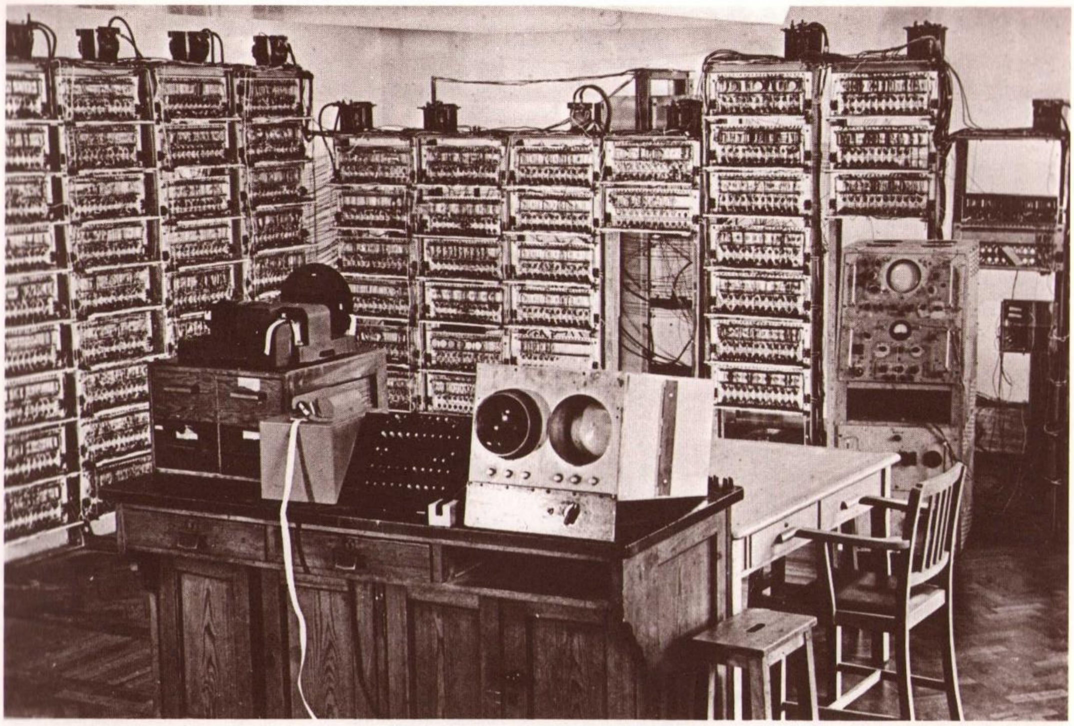

Meg, the prototype Manchester University 'Mark 2' computer, was first operational in May 1954. It was a serial/parallel design, incorporating hardware floating point arithmetic.

Distributed electromagnetic delay lines provided a cheap, flexible implementation for many of Meg's internal registers, etc (*reference 42*). A single drum backing store was provided for the prototype.

Co-operation with Ferranti led to the production Meg, known as the Ferranti Mercury (*reference 43*). This differed from Meg only in the details of the order-code repertoire and in the main store technology. Mercury's main store (*reference 44*) was 1024 × 40 bits of ferrite core, at 10 microsecond cycle time for a 10-bit short word. It was divided into four blocks for addressing purposes, consistent with the usage of short (10-bit), medium (20-bit) and long (40-bit) words. At that time the core store was considered to be more expensive than the corresponding CRT version, but more reliable and physically smaller. Mercury had a backing store consisting of four drums, each holding 4096 × 40-bit words and having a 17.28 milliseconds revolution time. Parity checking was provided on main and backing storage.

The first Mercury was delivered tó the Norwegian Defence Research Establishment in August 1957. The delay between design completion and commencement of Mercury sales

The Ferranti Mercury, the production version of Meg, was delivered to the first customer in August 1957. Mercury had a 10 microsecond cycle-time core store backed by four drums (seen in the four cabinets to the right of the above picture.)

was partly due to marketing uncertainties, although the slower Ferranti Pegasus computer was at that time enjoying a successful sales record. However, NRDC was effectively underwriting some of the Pegasus development. Pegasus had nickel magnetostrictive delay-line plus drum storage and a compact 'packaged' technology, and 38 Pegasus 1 and Pegasus 2 machines were delivered between 1956 and 1962. The second production Mercury was delivered to Manchester University late in 1957 and had passed its acceptance tests by 5th February 1958. The cost of the machine was met partly from Mark 1 earnings and from design know-how supplied to Ferranti and partly by a £48,000 grant from the University Grants Committee (UGC). It should be added that the Manchester Mercury was returned to the UGC free of charge in January 1963 upon commencement of the Atlas service (see below), and it was then installed by the UGC at Sheffield University. Nineteen Mercury machines were sold by Ferranti between 1957 and 1961, six of these going abroad. At least four Mercurys were still working in 1970. In 1958 the Mercury was held to be one of the most powerful computers available in Britain and good value for money. The market rival was the IBM 704. This had a fully parallel arithmetic unit giving a fixed-point add time (36 bits) of 24 microseconds, and a larger core store. Although a faster machine, it cost five times as much as Mercury.

Manchester University's Mercury provided a potential service to outside users in much the same way as the Mark 1 had done, except that by about 1961 the available time was almost all taken up by internal University departments. The rate charged to outside users was £50 per hour. Amongst the notable software developments on Mercury during 1958-62 was the high-level language Mercury Autocode (*reference 45*), first operational in basic form in

March 1958, with additional facilities for matrix manipulation, etc finished by August 1959. Brooker was assisted in this by a team which included B Richards, who had first used the Mark 1 as a research student working on morphogenesis. Brooker's work on phrase structure and formal language description (*reference 46*) was begun in about September 1958; this was initiated independently of the work elsewhere which is contained in the 1958 Algol report. The Manchester research led to the *Compiler Compiler* (*reference 47*), which was the first successful attempt at a formal language for the transparent description of syntax and semantics. This paved the way for much of the compiler writing effort on the Atlas project. D Morris, who joined the group in March 1960, assisted with the Compiler Compiler and subsequent system software. His first introduction to computing had come a few years earlier, whilst using the Mark 1 for civil engineering research.

The experimental transistor computer designed at Manchester University, first operational in November 1953. It is believed to be the first transistor computer to come into operation anywhere. The drum store, shown on the left of the picture, was originally used for Mark 1 in 1949

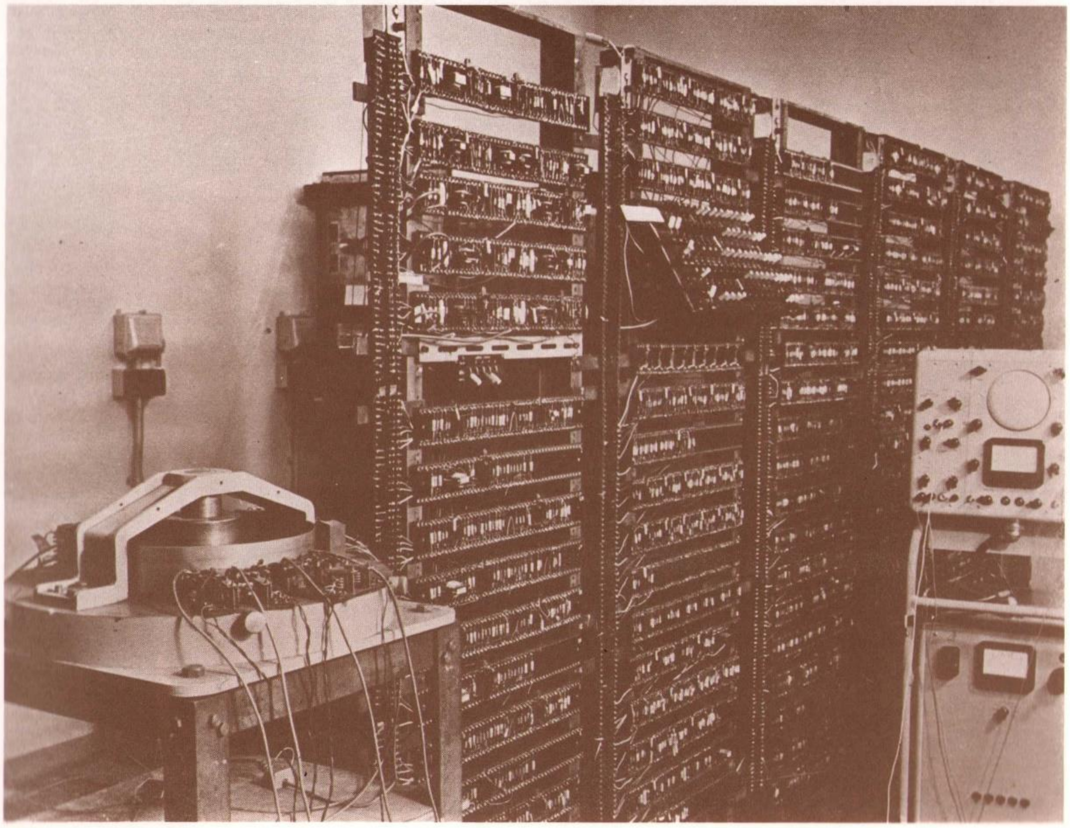

The transistor computer

In parallel with the Meg project it was decided to build a relatively small and economic computer. By the time the design got under way in 1952 it was clear that this could also provide valuable experience in the use of the recently-introduced transistors ('crystal triodes') as replacements for thermionic tubes. The transistors then available were germanium point-contact devices, which were less reliable than thermionic tubes but had the advantage of consuming far less power. Another advantage of the point-contact device was that a two-state flip-flop could be implemented using only one transistor (the current gain was greater than unity).

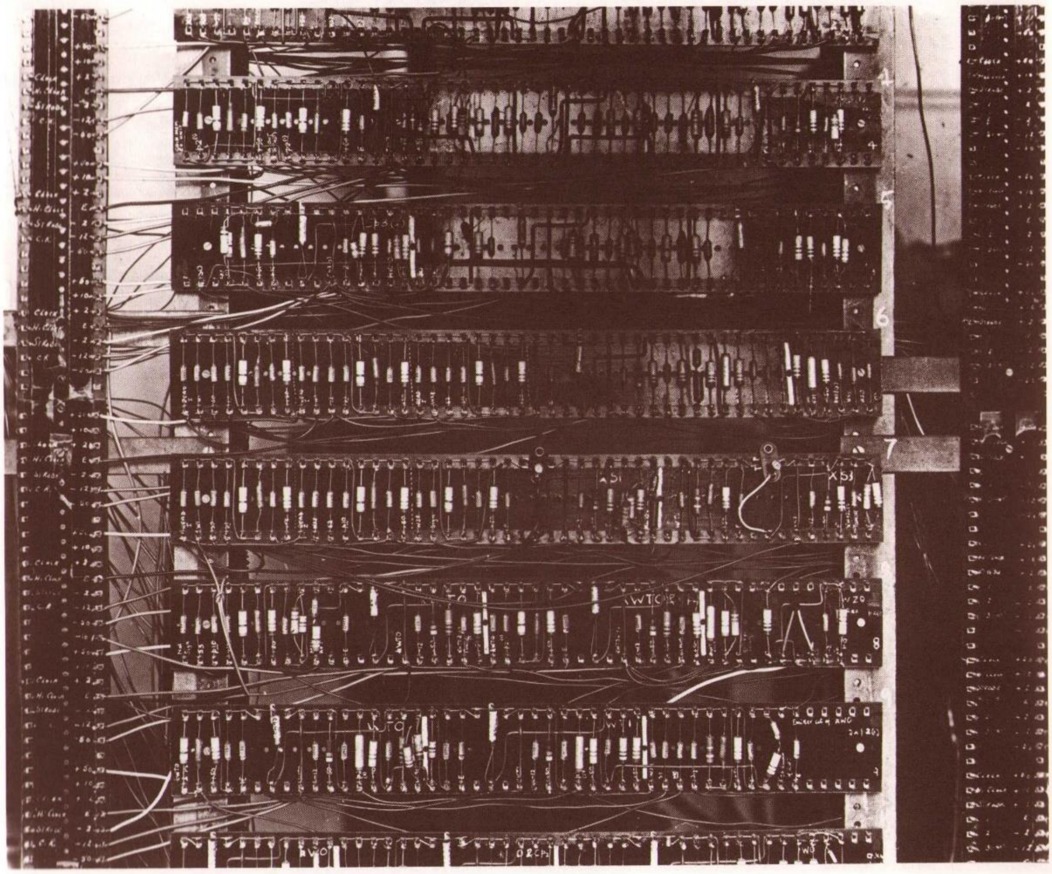

A close-up of some of the point-contact transistor circuits. For the production version, junction transistors mounted on printed-circuit boards were used

The Metropolitan-Vickers MV 950 based on the Manchester University transistor computer. The MV950 was completed in 1956 and six were produced.

Kilburn's computer design team within the Electrical Engineering Department in 1952 basically numbered five people including himself: D B G Edwards and G E Thomas working on Meg, and R L Grimsdale and D C Webb working on the transistor machine. G B B Chaplin and other research students in a different group within Electrical Engineering contributed to discussions of transistor circuit design (eg *references 48, 49*). E M Dunstan joined Kilburn's team in October 1954 at the end of the Meg development, to work on drums. From about 1952 right through until the start of the Atlas project there were usually about two Ferranti engineers attached to Kilburn's team for development and liaison purposes.

Two versions of the prototype transistor computer were commissioned in November 1953 and April 1955 respectively (*reference 50*). Both versions had a psuedo two-address (or 1 + 1) instruction format consistent with the fact that the only store was a single drum, with consequent latency problems. As with some other delay-line type storage computers, the address of the next instruction was contained within each instruction, thus facilitating 'optimum programming'. One track of the drum was used for the main accumulator and one track for the instruction register, in order to simplify the processor circuitry. The word length was 44 bits, divided into four 'syllables' for an instruction. The first computer (November 1953) had a simple seven-function order code and one track of 64 words for main storage. In the second computer (April 1955) the order code and storage were extended and a hardware multiplier included. One drum track formed an 8-word B-store which was also available as a set of 'fast' general registers. Arithmetic was serial, with a pulse rate of 125,000 per second. The instruction times were related to the 30 milliseconds drum revolution time, thus being slow compared to the Mark 1 performance.

The 1955 Machine had a total of 200 point-contact transistors and 1300 point-contact diodes and had a power consumption of 150 watts. Considerable reliability problems were experienced with early batches of transistors and the average error-free run in 1955 is quoted as only 1.5 hours. It is interesting to note that at first the computer's drum amplifiers used point-contact transistors. Later on in the project point-contact devices were replaced by junction transistors as soon as these became available, with a consequent improvement in stability, reliability and noise. Finally however, thermionic tube amplifiers were used for the drum since contemporary junction transistors did not quite have a good enough high frequency response. It may be inferred that perhaps the most important impact of this computer on the Manchester team was the experience it provided in transistor circuit techniques. The November 1953 machine is believed to have been the first transistor computer to come into operation anywhere.

The design of the experimental transistor computer was subsequently adopted by the Manchester firm of Metropolitan Vickers (later AEI and now GEC), who converted all the circuits to use the more reliable junction transistors. The production version, known as the MV950, was completed in 1956 and six machines were made. These computers operated successfully in various departments within Metropolitan Vickers for about five years, during which time they are said to have had impressive reliability records.

MUSE and Atlas

Aims and resources

By the mid-1950's Britain was falling behind the United States in the production of high-performance computers, as was confirmed in a report dated May 1956 to the DSIR Advisory Committee on High Speed Calculating Machines (The Brunt Committee). B W Pollard of Ferranti, speaking in April 1956, told a computer design conference that "there is in this country a range of medium-speed computers, and the only two machines which

are really fast are the Cambridge EDSAC 2 [*reference 52* – working by 1957] and the Manchester University Mark 2, although both are still very slow compared with the fastest American machines. We ought to be a little worried that no serious effort is being made to produce at least one really large fast machine'' (*reference 51*). By the autumn of 1956 Kilburn's team had initiated such an effort, known then as the Muse ('microsecond') computer.

Whilst Muse was in its early stages Kilburn had talked to several potential users of large computers including the Atomic Energy Authority (AEA) and some big commercial

M J Lanigan (left) and Tom Kilburn with part of the Muse computer in about 1959. The Muse project was subsequently developed into the Atlas computer by a joint University/Ferranti team under Kilburn organisations. From the former came the design specification for an instruction speed approaching one order per microsecond, whilst from the latter came the need to attach a large number of peripherals of various types. Both classes of user needed an immediate-access storage capacity far in excess of anything then available. It became clear at the outset that special techniques would be needed if efficient use of equipment and rapid turnround were to be achieved in a computer system consisting of many interconnected units, each of widely differing speeds. This applied not only to slow and fast peripherals and backing stores, but also to high-speed transistor logic circuits interfacing with medium-speed moderate-sized main stores. The special techniques that were eventually employed included what are now known as multiprogramming, job scheduling, spooling, interrupts, pipelining, interleaved storage, autonomous transfer units, virtual storage and paging – though none of these techniques had been invented when the project started in 1956.

It was clear to Kilburn that the Electrical Engineering Department at Manchester had nowhere near the resources to see through such a project. (It also seemed clear to some contemporary commentators that the lavish scale of Muse would be more than sufficient to satisfy the computing demands of the entire country and that the whole project was rather ridiculous!) Kilburn spent the next 18 months attempting to get financial suport for Muse, whilst initial design-work and assessment of technology was in progress. Over this period the Brunt Committee and Lord Halsbury of the NRDC were also trying to initiate a British high-speed computer project, spurred by reports of the massive scale of the Univac LARC and IBM STRETCH developments in the United States. During 1957 and 1958 many meetings were held between the NRDC and: industry (principally Ferranti and EMI); research groups (principally Manchester University and the Royal Radar Establishment [the successor to TRE]); and potential users (principally the Atomic Energy Authority). An attempt to analyse these complex discussions has been made elsewhere, to explain the delays which occurred (*reference 53*). It is sufficient to say that the factors influencing decisions were not easy to quantify accurately in 1957/58.

From Kilburn's viewpoint the plain facts were that Ferranti at first felt unable to support the Muse proposal and government finance was not forthcoming. The University team therefore went ahead with a limited version of Muse, with D B G Edwards responsible for hardware and R A Brooker for software, using normal departmental resources including the Mark 1 Computer Earnings Fund. However, in October 1958 (confirmed officially in January 1959) Ferranti decided to become involved in the project and in May 1959 NRDC provided £300,000 to Ferranti, repayable through a levy on eventual sales. Some idea of the financial uncertainty of the Muse involvement from Ferranti's viewpoint can be imagined from various estimates of the works cost of the first production model, which ranged from £375,000 in February 1959 to £650,000 (November 1960) to £930,000 (October 1962).

By the start of 1959 the computer had been re-named Atlas and was thereafter developed as a joint University/Ferranti venture under Tom Kilburn. At this point the hardware design team numbered about 13 University and 6 Ferranti people; the software numbers were about 7 (mostly on compilers) and 2 (on operating systems) respectively. Two years later when commissioning was in full swing the Ferranti hardware contribution had increased to about 18 engineers and the software contribution to about nine (including four test programmers). Production facilities and ancillary support at Ferranti's West Gorton factory are not included in these figures. It is not possible to mention all who contributed to Atlas. Amongst those most closely associated with the project under Kilburn were:

The Manchester University Atlas was commissioned in December 1962, at which time it was the most powerful computer in the world. This photograph shows two of the five central processor bays: the centre one contains the 8k fixed (read-only) store and the right-hand bay contains the magnetic tape autonomous transfer unit

Hardware/Architecture	*Compilers/Operating System*
D B G Edwards	R A Brooker
F H Sumner	D Howarth
D Aspinall	R B Payne
M J Lanigan	D Morris
E C Y Chen	I R MacCallum
G Haley	R H Kerr
E T Warburton	J S Rohl
K F Bowden	J Clegg
E M Dunstan	M T Wyld

An analysis of 15 Atlas hardware patent applications filed between October 1957 and February 1962 is shown in table 3.

By 1959 the hardware design of Atlas had been substantially completed (*reference 54*), plans for the compilers were in hand using the Compiler Compiler as a production tool, but the magnitude of problems associated with the Supervisor (operating system) were not

Table 3: An analysis of the inventors responsible for 15 Atlas patents, filed between October 1957 and February 1962 (with acknowledgement to NRDC)

name	dates at Manchester for computer design purposes	number of patents of which an inventor or co-inventor
T Kilburn	Dec '46 onwards	14
D B G Edwards	Oct '48 onwards	9
F H Sumner	Feb '58 onwards	2
M J Lanigan	Nov '57 – Dec '62	2
D Aspinall	Nov '57 – Sept '70	1
E C Y Chen	Ferranti	1

fully realised. It was with the Supervisor that Ferranti, in the form of David Howarth, made a fundamental contribution to the Atlas design. Howarth had recently joined Ferranti, having previously been an early user of Mark 1 whilst working at RRE. His capability and tremendous capacity for work – he had obtained a physics PhD before the age of 22 – led to the implementation of a Supervisor consisting of some 35,000 machine instructions with a team which at no time numbered more than six programmers. A lot of hard work by a comparatively small group of people characterised the whole Atlas development at Manchester during this period. The time schedules were made more acute by the planned shut-down of the Mercury computer on 31st December 1962 and the transfer of the user-service to the prototype Atlas – a commitment which was made before the magnitude of the Atlas commissioning task had been fully appreciated. To put this in perspective the LARC and STRETCH projects in the United States were also having problems and they were both to become 'failures' compared to Atlas.

Performance

Atlas was officially inaugurated on 7th December 1962 by Sir John Cockroft of the Atomic Energy Authority. It was considered to be the most powerful computer in the world, the Ferranti salesmen equating it to four IBM 7094's. There are some 25 published papers describing Atlas, of which references 55 to 59 are a representative sample. To summarise, Atlas had a 48-bit word, single address instruction format, allowing for double B-modification and a user virtual address space of one million words. There were 127 half-word B-registers, mostly held in a fast 0.7 microsecond cycle time core store. System routines and some frequently-used user library routines were held in an 8K read-only store having an access time of 0.3 microsecond. The core-plus-drum one level store used fixed-sized 512-word pages, 32 page address registers forming an associative store for virtual-to-real address translation. A 'drum learning program' held in read-only memory attempted to optimise page swapping by accumulating statistics on page utilisation. The Manchester Atlas had the following storage hierarchy:

main store: 16K*, 2 microsecond cycle time four-way interleaved. Backed by 4* drums, each 24K, 12 millisecond revolution time, capable of transferring one block of 512 words every 2 milliseconds.
system working store: 1K* 2 microsecond cycle time
system read-only store: 8K, 0.3 microsecond access
magnetic tapes: 8* decks (on 8 channels), each 90kc/s giving one word every 88 microseconds
file disc (added in 1967): 16M words
(*: increased on subsequent production Atlases)

The engineer's console and central processor of the Science Research Council's Atlas computer at Chilton (near Harwell). The Chilton Atlas (1964) was the third Atlas to be sold. Ferranti also produced three smaller versions generally known as Atlas 2's

The interrupt structure allowed for up to 512 peripherals. The Manchester Atlas had about 17 conventional input/output devices. Two high-speed data links, an on-line X-ray crystallographic unit and an experimental speech input/output device were added later. An idea of the instruction speed may be obtained from the following orders, measured on 7th December 1962:

fixed-point B addition	1.59 microseconds
floating-point add, no modification	1.61 microseconds
floating-point add, double modification	2.61 microseconds
floating-point multiply, double modification	4.97 microseconds
floating-point division:	10.66 microseconds min to 29.80 microseconds max

A primitive user-service was available in 1963 using a Supervisor without multiprogramming capability and Atlas Autocode and Mercury Autocode as the only high-level languages. The full Supervisor became available from January 1964 followed by other languages such as Algol and Fortran. By the mid-1960's, when the useful time was shared between the Ferranti and University computing services, anything up to 1000 programs were being run in a 20 hour 'day'. It was part of an agreement with Ferranti that a percentage of the Ferranti income obtained from outside users should be paid to the University Computer Earnings Fund. Ferranti charged up to about £500 per hour for Atlas time. By 1969 the annual value of computing to the University, including the Institute of

An overall view of the data-preparation area for the Chilton Atlas computer. Through the glass viewing window in the background may be seen the main input/output peripheral equipment

Science and Technology, was estimated at £720,000 if costed on the open market. The Manchester Atlas was also used by five remote users via high-speed data links – a pioneering venture in Britain. These remote sites were: Jodrell Bank radio telescope, Edinburgh University, Nottingham University, Ferranti's London Bureau and The Nuclear Power Group. To the present-day user Atlas would seem to have a very limited concept of file manipulation and no access via time-sharing terminals; in all other respects however the facilities and throughput offered to the user would probably be equivalent to those obtained from typical large computers designed ten years after Atlas. The Manchester Atlas was closed down with appropriate valedictory libation on 30th September 1971.

Ferranti sold two full versions of Atlas – one to a joint London University/British Petroleum consortium and the other to the National Institute for Research into Nuclear Science (NIRNS) at Harwell. Hardware for the former was delivered in 1963 and for the latter in December 1964. The Harwell site was subsequently re-constituted as the Science Research Council's Atlas Computer Laboratory at Chilton – which in fact lies a few yards outside the boundary fence of the Atomic Energy Authority's Harwell establishment. The Chilton (Harwell) Atlas was the largest Atlas installation, having 48K of main core and 32 magnetic tape decks. This machine was made available to all British Universities until its shut-down in March 1974.

Subsequent developments

In February 1962 Ferranti provided Cambridge University with some units of Atlas hardware on special terms, in return for aid in developing a simpler version of Atlas. This arrangement suited Cambridge because it provided a test-bed for incorporating their own hardware and software ideas which would hopefully result in a reasonably powerful service computer for the University. The product of this was the Cambridge Titan machine (*reference 60*), first operational in the summer of 1963. It lacked the one-level store and system read-only store of an Atlas 1 but incorporated an experimental 32-word slave store

The input/output peripherals for the Chilton Atlas. The central processor lies on a lower floor beneath this room. Many Atlas ideas, including virtual storage, are used in today's large computers

The initial design for Manchester University's fifth computer, MU5, was begun in 1966. The photograph shows a protoype 16-bit adder, using emitter-coupled logic, in about 1968

(which however was never fully utilised due to persistent hardware faults). The main store size was originally 32k at 6 microseconds cycle time, later increased to 128k at 2.5 microseconds. Cambridge wrote a special Titan operating system with interesting file manipulation and time-sharing features. Ferranti (later International Computers and Tabulators) sold two production Atlas 2's, which were essentially 128K Titans without the slave store. One of these went to the Atomic Weapons Research Establishment at Aldermaston in 1963 and the other, after a spell at the West Gorton factory, went to the government-sponsored Computer Aided Design Centre near Cambridge in 1966 – where it is still working.

In the early 1960's Ferranti's only competitor in the large machine market was Control Data Corporation. CDC have said that information presented in Kilburn's Muse paper in

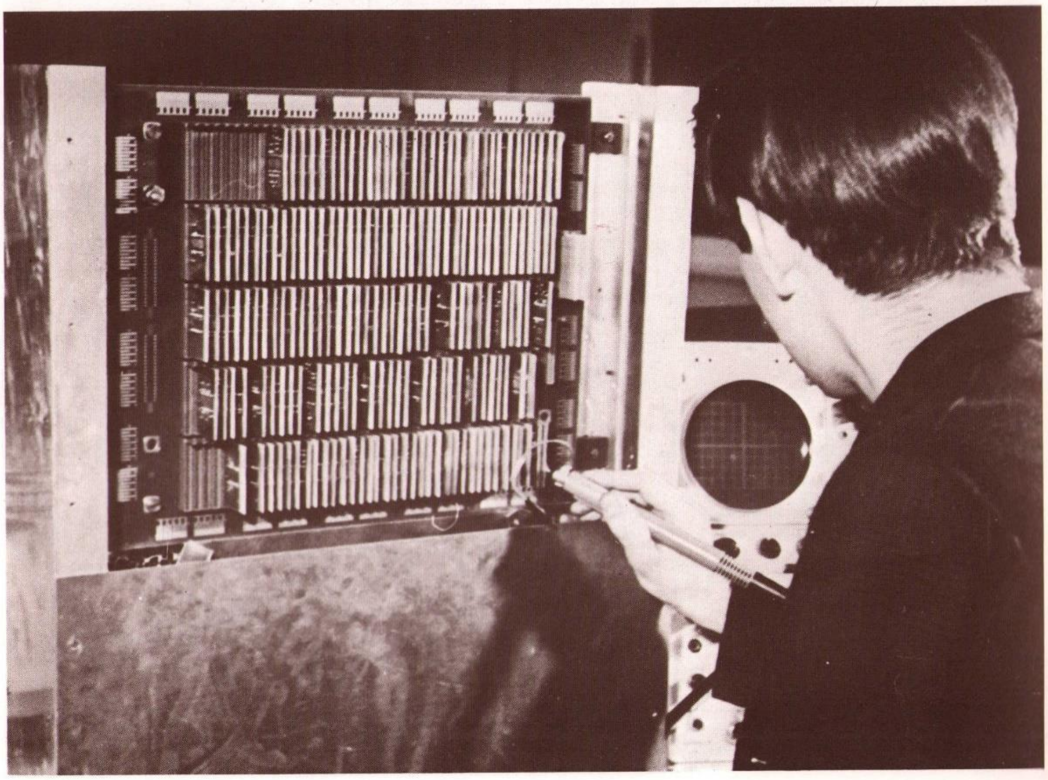

1959 (*reference 54*) helped them significantly in developing the CDC 6600 computer sooner than originally expected. In the event CDC were able to secure orders such as the Australian CSIRO installation, for which Atlas was originally a prime contender. In 1963 Ferranti's main computer interests were taken over by ICT (later International Computers Limited) and a decision was made to go for a broader marketing base with the 1900 series machines as the principal line. It is disappointing that circumstances prevented more Atlases from being sold. There is no doubt that the main impact of Atlas in the world-wide computer situation has been the wealth of new ideas which the project introduced – ideas, such as virtual storage, which are finding a wide use in present-day large machines

MU5

In 1964 the computer group within Electrical Engineering, consisting of 12 staff, formed a separate Department of Computer Science under Tom Kilburn as Professor. The first undergraduates were accepted into a three years honours degree course in October 1965. The present staff numbers 32 including four Professors (Kilburn, D B G Edwards, F H Sumner and D Morris).

Preliminary planning for a new large computer project was begun in the Department of Computer Science in about 1966, with proposals for system architecture and the development of an integrated associative storage technology (*references 61, 62*). The project, which was later christened MU5, was envisaged as having a speed about twenty times that of Atlas. One of the main design aims was to produce an architecture capable of running high-level language programs efficiently – in contrast to the significant discrepancy

Part of the MU5 computer room in 1974. The MU5 project was funded by the Science Research Council, and ICL provided manufacturing facilities

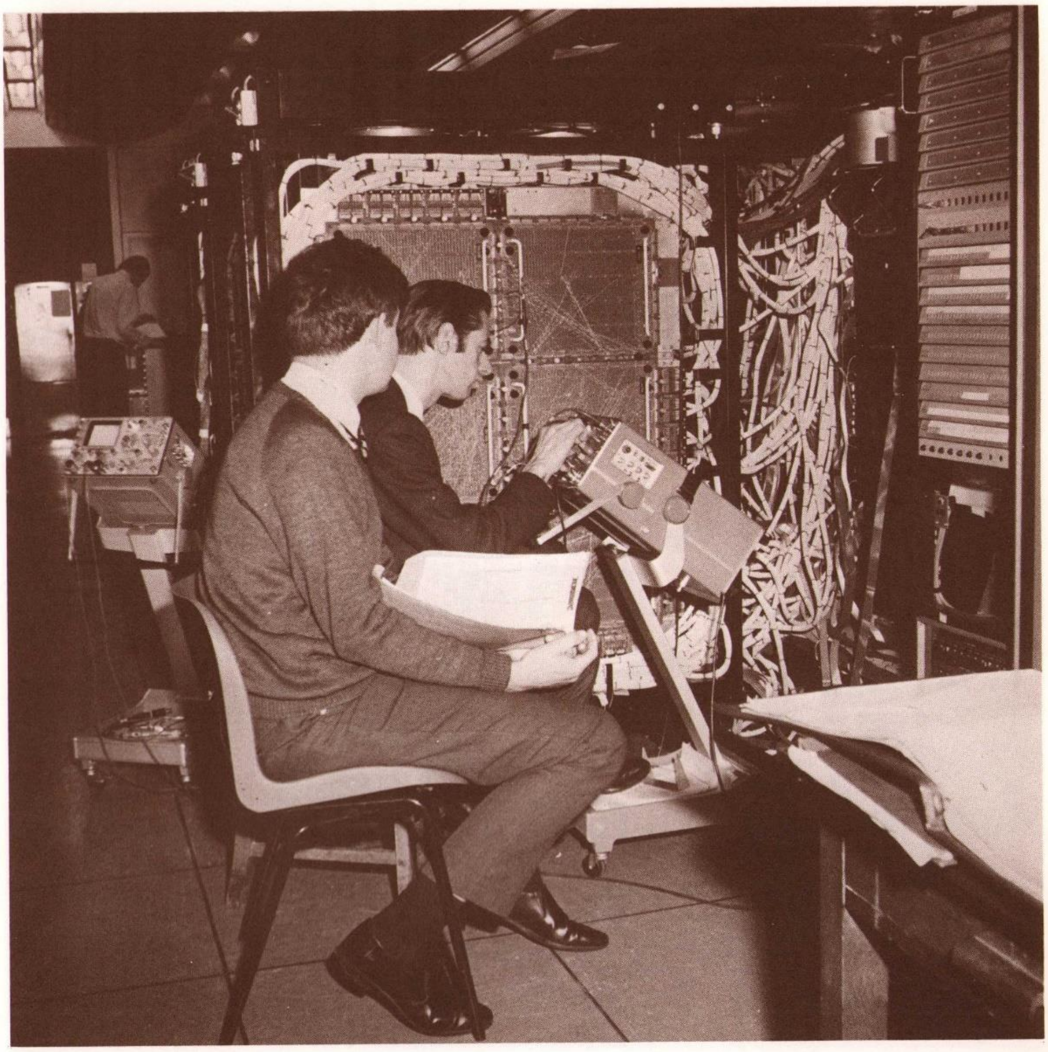

observed between hand coding and code generated by compilers for existing conventional computers. This was to be achieved not only by careful formulation of the order code but also by the use of hardware such as associatively-accessed *name stores* to give the effect of optimised central registers. A segmented virtual store with variable-sized pages was planned, facilitating the sharing and protection of resources necessary for a real-time multiprogramming environment.

In 1968 the Science Research Council (SRC) awarded a five-year grant of £630,446 for the project and ICL agreed to give significant help with production facilities. Part of the SRC grant purchased an ICL 1905E computer which was used between early 1969 and early 1972 for detailed MU5 logic simulation and as a test-bed for an asociative store. The associative store formed a paging unit which gave the 1905E a virtual address structure similar to that of the eventual MU5. This allowed pilot versions of the MU5 system software to be developed from about 1969 onwards, in parallel with the MU5 hardware. The paged 1905E is now linked into the MU5 multicomputer system.

The initial design contributions for the MU5 processor came from the following nucleus of Computer Science Department staff: Tom Kilburn, D B G Edwards, D Aspinall, D Morris, J S Rohl and F H Sumner. Software considerations had a dominant influence on discussions from the early stages, and for the high speeds that were planned sophisticated integrated circuits and interconnection technology were obviously vital. Because of the size of the project collaboration was established with the ICL team at West Gorton, who had been working for some time on advanced technology. By 1971 the MU5 design team numbered 16 Computer Science staff backed by some 25 research students, together with about 19 engineers from ICL. The parts were manufactured by ICL at West Gorton using the new multilayer board plant set up for their 1906A/S machines and 2900 project. At the same time ICL was working on the design of a new computer range announced in October 1974 as the 2900 series. The architecture of the 2900 series owes much to and has a great deal in common with MU5. By October 1974 MU5 was running ALGOL and FORTRAN benchmarks and monitoring of the hardware and software performance had begun.

The MU5 project is described in over 40 published papers, of which references 61 to 68 are a sample. MU5 is not yet 'history' so no overall evaluation will be attempted here. Perhaps in a few years there will be an MU6. If so this book will eventually be updated, continuing the unique story of Manchester computer design.

MU5 during commissioning: a logic door has been opened to allow access to signals on the multi-layer printed-circuit platters. The particular group of circuits in the photograph include an associatively-accessed fast operand buffer or 'Name Store'

References

1. Michie, D: *The Bletchley machines*, unpublished memorandum, June 1972, reprinted in *The Origins of Digital Computers: Selected papers*, edited by Brian Randell (Springer-Verlag 1973)
2. Goldstine, H H and Goldstine, A: *The Electronic Numerical Integrator and Computer (ENIAC)*, MTAC, Vol 2, No 15, 1946, 97-110, reprinted in Randell, (*see ref 1*)
3. Wilkes, M V: *Babbage as a Computer Pioneer*, Proceedings of the Babbage Memorial Meeting, held Oct 1971 (British Computer Society, 1972)
4. *A Discussion on Computing Machines*, organised by the Royal Society on 4th March 1948 with Prof Newman in the chair. His contribution, entitled *General Principles of the Design of All-Purpose Computing Machines*, is reproduced in Proc Royal Soc, Vol 195A, 1948, 271-74
5. Royal Society Minutes of Council, 1945-48, Vol 17
6. Newman, M H A: *A Status Report on the Royal Society Computing Machine Laboratory*, prepared for an internal committee of Manchester University, 15th October 1948
7. Burks, A W, Goldstine, H H, and Von Neumann, J: *Preliminary discussion of the logical design of an electronic computing instrument*, Part 1, Vol 1, Princeton: Institute for Advanced Study, June 1946 (Reprinted, together with the first draft of a report on the EDVAC dated June 1945, in Randell – *see ref 1*)
8. *Biographical Memoirs of Fellows of the Royal Society*, Vol 4, 1958 : Hartree, D R, 103ff
9. Hartree, D R: *A historical survey of digital computing machines* Proc Royal Soc, Vol 195A, 1948, 265-71 (*see also ref 4*)
10 *The Telecommunications Research Establishment Parallel Electronic Digital Computer*, described in Chapter 10 of *Faster than Thought*, ed B V Bowden (Pitman, 1953)
11. *Report of a conference on high speed automatic calculating machines*, Cambridge, 22-25 June 1949 (University Mathematical Laboratory, Cambridge, 1950)
12. *Faster Than Thought: A symposium on digital computing machines*, ed B V Bowden (Pitman, 1953)
13. Turing, A M: *On computable numbers, with an application to the Entscheidungsproblem*, Proc London Math Soc, Ser 2, Vol 42, 1937, 230-65. Corrigenda in Vol 43, 544-46
14. *Contribution to a discussion at the Cambridge 1949 conference* (*see ref 11*, 133)
15. Kilburn, T: *A storage system for use with binary digital computing machines*, PhD thesis, University of Manchester, 1948
16. Williams, F C and Kilburn, T: *A storage system for use with binary digital computing machines*, Proc IEE, Vol 96, part 2, No 30, 1949, 183ff
17. Williams, F C, Kilburn, T, and Tootill, G C: *Universal high-speed digital computers: A small-scale experimental machine*, Proc IEE, Vol 98, Part 2, No 61, Feb 1951, 13-28
18. Williams, F C and Kilburn, T: *Electronic digital computers*, letter to Nature, Vol 162, Sept 1948, 487
19. Williams, F C: *Early computers at Manchester University, (to be published in the Radio and Electronic Engineer)*
20. *Turing, Sara: Alan M Turing*, (Heffer, Cambridge, 1959)
21. Kilburn, T: *The Manchester University digital computing machine*, 119-22 in the Cambridge 1949 Conference (*see ref 11*)
22. Kilburn, T: *The University of Manchester Universal High Speed Digital Computing Machine*, Nature, Vol 164, Oct 1949, 684-91
23. Thomas, G E: *Magnetic storage*, 75-80 in the Cambridge 1949 conference (*see ref 11*)
24. West, J C and Williams, F C: *The position synchronisation of a rotating drum*, Proc IEE, Vol 98, Part 2, No 61, 1951
25. Williams, F C and Kilburn, T: *The University of Manchester computing machine, Inaugural Conference of the Manchester University computer, July 1951, 5-11. This paper was also presented at the Joint AIEE-IRE Computer Conference, Philadelphia, Dec 1951*
26a. *Kilburn, T: The new computing machine at the University of Manchester*, Nature, Vol 168, 1951, 95-6
26b. Williams, F C, Robinson, A A, and Kilburn, T: *Universal high speed digital computers: serial computing circuits*, Proc IEE, Vol 99, Part 2, No 68, April 1952
26c. Edwards, D B G: *The Manchester University high speed digital computer*, Journal Brit IRE, April 1954, 269-78
27. Crawley, H J: *A summary of the activities of the National Research Development Corporation in connection with electronic computer development*, NRDC paper 132, Feb 1957
28. Tootill, G C: *Informal report on the design of the Ferranti Mark 1 computing machine*, Ferranti internal memorandum, Nov 22nd 1949
29a. Pollard, B W: *The design, construction, and performance of a large-scale general-purpose digital computer*, Joint AIEE-IRE Computer Conference, Philadelphia, Dec 1951
29b. Pollard, B W and Lonsdale, K: *The construction and operation of the Manchester University Computer*, Proc IEE, Vol 100, Part 2, 1953, 501-12. For comments on reliability see also 33,34 of the Inaugural Conference report (*see ref 25*)
30. Somers, C (the Ferranti archivist): internal memorandum
31. Robinson, A A: *Multiplication in the Manchester University high-speed digital computer*, Electronic Engineering, Jan 1953, 6-10
32. Woodger, M: *A comparison of one and three address codes*, Manchester Inaugural Conference, 1951, 19-23 (*see also ref 25*)
33. Turing, A M: *Generation of random numbers*, Appendix 2 in the Ferranti internal memorandum (*see ref 28*), Nov 22nd, 1949

34. See under *Electronic Computers* (190-191) in the 64th edition of Pears Cyclopaedia (London, 1955)
35. Kilburn, T, Tootill, G C, Edwards, D B G, and Pollard, B W: *Digital computers at Manchester University*, Proc IEE, Vol 100, Part 2, 1953, 487-500
36. Williams, F C: *IEE convention on digital computers – Introductory lecture*, Proc IEE, Vol 103, Part B, Supp 1-3, 1956, 3-9
37. Williams, F C: Internal report on the computing machine running costs, University of Manchester, June 1955
38. Newman, M H A: *Obituary notice of A M Turing*, Biographical memoirs of Fellows of the Royal Society, Vol 1, 1955, 253-63
39. Brooker, R A: *An attempt to simplify coding for the Manchester electronic computer*, British Journal of Applied Physics, Vol 6, Sept 1955, 307-11
40. Brooker, R A: *The programming strategy used with the Manchester University Mark 1 computer*, Proc IEE, Vol 103, Part B, Supp 1-3, 1956, 151-57
41a. Kilburn, T, Edwards, D B G, and Thomas, G E: *The Manchester University Mark 2 digital computing machine*, Proc IEE, Vol 103, Part B, Supp 1-3, 1956, 247-68
41b. Edwards, D B G: *The design and operation of a parallel type cathode ray tube storage system*, Proc IEE, Vol 103, Part B, Supp 1-3, 1956, 319-26
41c. Williams, F C, Kilburn, T, Witting, C N W, and Hoffman, G R: *Recent advances in cathode ray tube storage*, Proc IEE, Vol 100, Part 2, No 77, Oct 1953, 523-43
42. Thomas, G E: *The use of electromagnetic delay lines in the Manchester University Mark 2 digital computing machine*, Proc IEE, Vol 103, Part B, Supp 1-3, 1956, 483-90
43. Lonsdale, K, and Warburton, E T: *Mercury: a high-speed digital computer*, Proc IEE, Vol 103, Part B, Supp 1-3, 483-90
44. Robinson, A A, Newhouse, V L, Friedman, M J, Bindon, D G, and Carter, I P V: *A digital store using a magnetic core matrix*, Proc IEE, Vol 103, Part B, Supp 1-3, 1956, 295-301
45a. Brooker, R A: *The Autocode programmes developed for the Manchester University Computers*, Computer Journal, Vol 1, 1958, 15-21
45b. Brooker, R A: *Further Autocode facilities for the Manchester (Mercury) Computer*, Computer Journal, Vol 1, 1958, 124-27
46. Brooker, R A, and Morris, D: *An assembly program for a phrase structure language*, Computer Journal, Vol 3, 1960, 168-74
47. Brooker, R A, MacCallum, I R, Morris, D, and Rohl, J S: *The Compiler Compiler*, Annual Review in Automatic Programming, Vol 3, 1963, 229ff
48. Chaplin, G B B: *The transistor regenerative amplifier as a computer element*, Proc IEE, Vol 101, Part 3, 1954, 298ff
49. Foster, J R: *A transistor digital computer: track selection systems for use with the magnetic drum store, and new computer circuits using junction transistors*, MSc thesis, University of Manchester, October 1954
50. Kilburn, T, Grimsdale, R L, and Webb, D C: *A transistor digital computer with a magnetic drum store*, Proc IEE, Vol 103, Part B, Supp 1-3, 1956, 390-406
51. Pollard, B W: *Contribution to a conference discussion*, reproduced in Proc IEE, Vol 103, Part B, Supp 1-3, 1956, 204
52. Renwick, W: *EDSAC 2*, Proc IEE, Vol 103, Part B, Supp 1-3, 1956, 277-79
53. Drath, P: *The relationship between science and technology: University research and the computer industry 1945-1962*, PhD thesis, University of Manchester, October 1973 (Concentrates on Manchester and Ferranti, 1948-62. Shaky on the computers but good on the economics and politics)
54. Kilburn, T: *Muse*, published in Information Processing (Butterworths, London, 1960) 433 (presented at the International Conference on Information Processing, UNESCO, Paris, June, 1959. Presentations on LARC and STRETCH were given in the same session)
55. Kilburn, T, Edwards, D B G, Lanigan, M J, and Sumner, F H: *One level storage system*, IRE Trans on Electronic Computers, Vol EC-11, No 2, April 1962, 223-35
56. Atlas storage details
a) Kilburn, T, and Grimsdale, R L: *A digital computer store with very short read time*, Proc IEE, Vol 107, Part B, No 36, Nov 1960, 567-72
b) Edwards, D B G, Lanigan, M J, and Kilburn, T: *Ferrite core memory systems with rapid cycle times*, Proc IEE, Vol 107, Part B, No 36, Nov 1960, 585-98
c) Edwards, D B G, Aspinall, D, and Lane, R: *Addressing magnetic tape on the Atlas 1 computing machine*, Computer Journal, Vol 8, No 4, Jan 1966, 323-29
57. Atlas arithmetic details:
a) Kilburn, T, Edwards, D B G, and Aspinall, D: *A parallel arithmetic unit using a saturated transistor fast-carry circuit*, Proc IEE, Vol 107, Part B, No 36, Nov 1960, 573-84
b) Sumner, F H, Haley, G, and Chen, E C Y: *The central control unit of the Atlas computer*, Proc IFIP Congress, 1962, 657-63
58. Atlas Supervisor details:
a) Kilburn, T, Howarth, D J, Payne, R B, and Sumner, F H: *The Manchester University Atlas operating system: Part 1, Internal organisation*, Computer Journal, Vol 4, No 3, 1961, 222-25. *Part 2, User's description*, Computer Journal, Vol 4, No 3, 1961, 226-29
b) Morris, D, Sumner, F H, and Wyld, M T: *An appraisal of the Atlas Supervisor*, Proc ACM National Meeting, 1967, 67-75
59. Atlas Compiler details:
a) Morris, D, and Rohl, J S: *The Atlas compiler*

system, Computer Journal, Vol 10, 1967, 227-30
b) Brooker, R A, Rohl, J S, and Clark, S R: *The main features of the Atlas Autocode*, Computer Journal, Vol 8, 1965, 303-10
60. Barron, D W, Fraser, A G, Hartley, D F, Landy, B, and Needham, R M: *File handling at Cambridge University*, Proc AFIPS Spring Joint Computer Conference, 1967, 163-67
61. Kilburn, T, Morris, D, Rohl, J S, and Sumner, F H: *A system design proposal*, Proc IFIP Congress 1968, Section D, 76-80
62. Aspinall, D, Kinniment, D J, and Edwards, D B G: *An integrated associative memory matrix*, Proc IFIP Congress 1968, Section D, 86-90
63. Kinniment, D J, and Edwards, D B G: *Circuit technology in a large computer system*, Proc IERE, Joint conference on computers – systems and technology, (London), Oct 1972, 441-50
64. Ibbett, R N: *The MU5 instruction pipeline*, Computer Journal, Vol 15, No 1, Feb 1972, 43-50
65. Woods, J V, and Sumner, F H: *Operand accessing in a pipelined computer system*, IEE Joint conference on computers – systems and technology, Nov 1974, 236-43
66. Standeven, J, Lanyado, S, and Edwards, D B G: *The MU5 secondary operand unit*, Proc IERE, Joint conference on computers – systems and technology, (London), Oct 1972, 429-40
67. Lavington, S H, Thomas, G, and Edwards, D B G: *The MU5 Exchange*, IEE Joint conference on computers – systems and technology, Nov 1974, 219-25
68. Morris, D, Detlefsen, G D, Frank, G R, and Sweeney, T J: *The structure of the MU5 operating system*, Computer Journal, Vol 15, No 2, 1972, 113-16
69. Capon, P C, Morris, D, Rohl, J S, and Wilson, I R: *The MU5 compiler target language and autocode*, Computer Journal, Vol 15, No 2, 1972, 109-12

about the author

Simon Lavington graduated in 1962 in Electrical Engineering at Manchester University, where he has worked ever since. He gained his MSc (1964) and PhD (1968) for hardware and software research into automatic speech recognition on the Atlas computer. Since 1968 he has worked on the MU5 multi-computer project. He has been a senior lecturer in computer science since 1974 and was until recently a British Computer Society examiner, with special responsibility for syllabus redesign. He is a United Nations consultant on hardware systems courses at the International Computer Education Centre, Budapest, and has been the author of a television series and accompanying book on the *Logical Design of Computers*.

about this book
The world's first stored-program computer was only the beginning of the continuous activity of computer design at the University of Manchester since 1946. Five original computers have led, through collaboration with local industry, to five commercially-available derivatives. This book presents the story, in words and pictures, of these five computers and their relationship to contemporary machines produced elsewhere.

UK net price

50p

ISBN 0 85012 155 8