

Systems Architecture of the Ferranti Mark I and Mark I* computers.

Contents.

1. Historical development.
2. The early Manchester University computers, 1948 – 1949.
3. Overview of the Ferranti Mark I.
4. Summary of the main differences between the Ferranti Mark I and the Mark I*.

1. Historical development.

These Ferranti machines were the fully-engineered production versions of prototype computers designed and built at the University of Manchester. The first of the prototypes was called the *Baby*, or more formally the Small Scale Experimental Machine (SSEM). The *Baby* first ran a program on the morning of Monday 21st June 1948, thereby being the world's first electronic stored-program computer to become operational. The SSEM was a minimal computer, primarily intended to test a novel storage system that had been developed by F C Williams and T Kilburn, first at the Telecommunications Research Establishment and then, from January 1947 at the Electro-technical Department of the University of Manchester where Professor Freddie Williams moved to become head of Department. Tom Kilburn was seconded to Manchester, to help with the storage research.

The novel storage system, later to be called the *Williams-Kilburn tube*, depended upon regenerating a pattern of electrostatically- charged spots on the inner surface of a cathode ray tube. The system is thus sometimes called CRT storage, or electrostatic storage. Each charged spot represented a binary one or zero. Two varieties of spots were tried: dot-dash, and focus-defocus, the latter being used on the Ferranti computers. The system was patented and eventually used under licence on several early computers, including the IBM 701 and 702. Unlike the sequential mercury delay lines of the time, the CRT system had the advantage that access-time was independent of address. That is, the CRT system was *random-access*, in the same sense as a modern RAM.

The 1948 SSEM had a single CRT storing thirty two 32-bit words and having a minimal instruction set of only seven operations. During the 12 months following its first operation, the small computer was enhanced in several respects: the word length was increased to 40 bits, the instruction set increased to 26 operations, a hardware multiplier was added, two index registers (known at Manchester as B lines) were provided, the primary memory was increased to 128 40-bit words and a drum backing store of capacity 1024 words was provided. Input and output was via 5-track teleprinter paper tape equipment. This enhanced computer was referred to as the *Manchester University Mark I* and sometimes as MADM (*Manchester Automatic Digital Machine*). It performed useful work during 1949/50, mostly investigating problems in mathematics. The design was passed to Ferranti at the end of 1949 and MADM itself was closed down in the summer of 1950 to await the arrival of the production version.

As a preliminary to describing the Ferranti Mark I, brief details of the Manchester University *Baby* (as at June 1948) and the Manchester University Mark I computer (as at October 1949) are now presented.

2. The early Manchester University computers, 1948 – 1949.

Figure 2.1 shows the layout of the Small Scale Experimental Machine, or *Baby*, in June 1948.

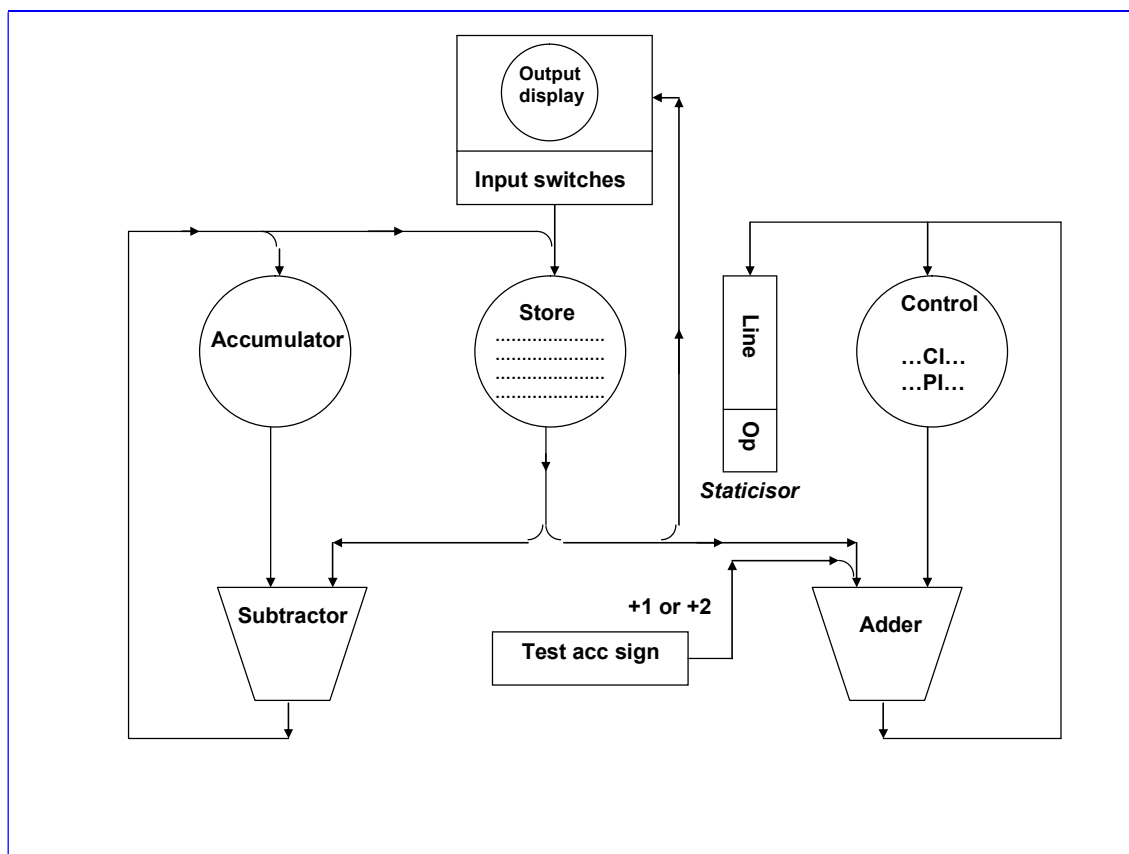


Figure 2.1. System diagram for the Manchester *Baby* (SSEM).

There are three Williams/Kilburn CRT storage units in Figure 2.1, each denoted by a circle. The one labelled *accumulator* holds a one-word (32-bit) line. The one labelled *Store* holds thirty-two lines of 23 bits each. The one labelled *Control* holds two 32-bit lines, used respectively for the Present Instruction PI and the Program Counter, CI. Instructions are 32 bits long, arranged as follows:

| 13 bits | | 3 bits | | 16 bits | |
|------------------------|----|----------------|----|-------------------|----|
| <i>Operand address</i> | | <i>Op code</i> | | <i>Unassigned</i> | |
| 0 | 12 | 13 | 15 | 16 | 31 |
| <i>Least-sig.</i> | | | | <i>most-sig.</i> | |

Only seven instructions (op codes) were available:

| Op code | Original notation | Modern description |
|--------------------|------------------------------|--|
| 0 | s to C | Absolute indirect unconditional jump |
| 1 | s + c to C | Relative indirect unconditional jump |
| 2 | - s to A | Load negative |
| 3 | a to S | Store accumulator |
| 4 | a - s to A | Subtract |
| 5 | - | <i>Not used (treated same as subtract)</i> |
| 6 | Test | Skip next instruction if Acc is negative |
| 7 | Stop | Halt |

The upper-case letters in the second column stand for the following:

S indicates the currently-addressed location in the store;

A refers to the accumulator;

C refers to Control (the *program counter*).

Lower case letters refer to the contents of the corresponding address (or unit).

By November 1948 plans were made to increase the word-length to 40 bits, with two 20-bit instructions in each word. There were to be two index, or modifier, registers known as B0 and B1. There was to be a double-length accumulator A, whose upper/lower halves were respectively called Am and Al. There was also a multiplier register R and a multiplicand register D – though the hardware multiplier did not exist at that stage. The main storage consisting of Williams/Kilburn tubes was to be 128 words (each of 40 bits). The instruction set, as defined in November 1948, contained 27 functions (operations) including facilities for double-length arithmetic. Logical AND, OR and NEQ (exclusive OR) were also provided. Overall, the instruction set reflected the intentions of the mathematicians at Manchester University to use the computer for problems in pure mathematics such as investigation of Mersenne primes and the Reimann Zeta hypothesis – see [Ref. 25 in section F1/X5.

By April 1949 the above machine was operational with its hardware multiplier. A 1024-word magnetic drum had also been attached as backing store though, at this stage, all drum transfers and all input/output from/to the 5-track teleprinter equipment was carried out by manual switches rather than by program. Nevertheless, some useful user-programs were run during the summer of 1949. Programmed input/output and drum transfers had been implemented by October 1949. By the end of 1949 a full specification of the instruction set for the production version of the university machine had been passed to Ferranti Ltd., who had been given a government contract to build a computer “to the specification of Professor Williams”.

The Manchester University Mark I (MADM) was closed down in August 1950, to await the arrival of the re-engineered Ferranti version. This also marked the point at which the academics, led by Tom Kilburn, started thinking of a successor computer which was to be ten times faster. The project was called *Meg*. This was in due course to be developed into the Ferranti Mercury. For more on the Manchester University computers, see references 20 and 25 in section F1X5.

3. Overview of the Ferranti Mark I.

The first Ferranti Mark I arrived at University of Manchester on 12th February 1951, thereby being the world's first commercially-available computer to have been delivered. Its design closely followed the 1949 university prototype. Confusingly, it was at first called the *Manchester Electronic Computer Mark II*. It was not until 1952, once the marketing possibilities had become clear, that the name was changed to the Ferranti Mark I. Two of these computers were delivered to customers – see list in F1X1. The design was then modified in minor respects (see later) and the new version was called the Ferranti Mark I* (pronounced *Mark One Star*). Seven of the updated versions were sold.

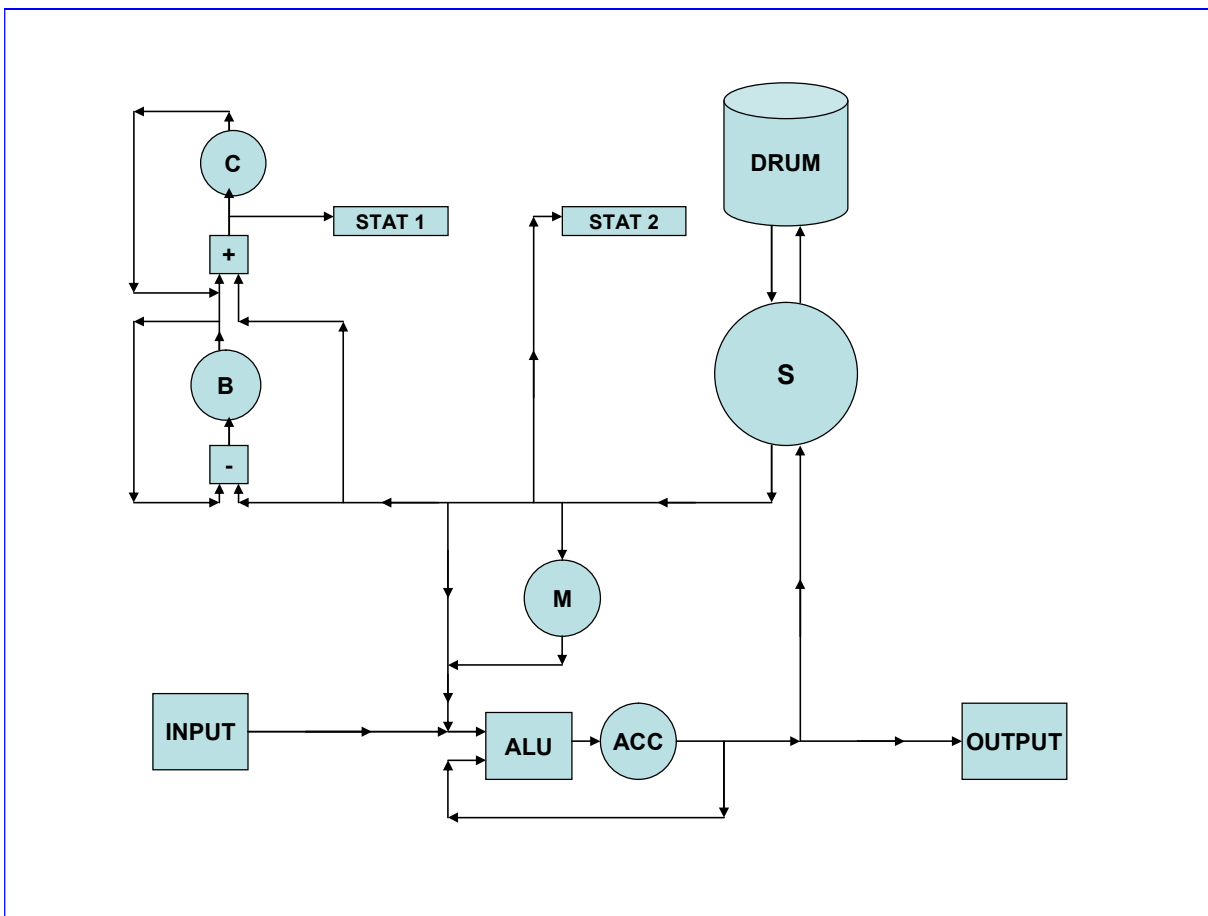


Figure 2.2. System diagram of the Ferranti Mark I.

In the above schematic, the circles represent Williams-Kilburn storage tubes which are used for the following purposes:

- S** primary memory, as eight tubes each storing 32 40-bit words making a total of 256 words. The bits on each tube were actually arranged as 64 lines of 20-bit half-words, addressable to the 20-bit half-word level.
- B** a single tube holding eight 20-bit index registers (called *B-lines*). This tube had its own subtractor, shown in the diagram as a box labelled '-'. B0, by convention, is normally arranged by the programmer to contain zero.
- ACC** the main double-length accumulator, arranged as four 20-bit lines.
- M** four 20-bit lines, holding a 40-bit multiplicand D and a 40-bit multiplier, R.

- C the control tube, storing the 20-bit present instruction PI and a 20-bit program counter (normally referred to simply as C). The C tube had its own adder, shown in the diagram as a box labelled '+'.

There were two flip-flop registers or *staticisors*, shown as thin rectangles in Figure 2.1. STAT 1 held a copy of the 20-bit present instruction. STAT 2 held the 20-bit control word used for programmed drum transfers, as described later. The drum held about 8K of 20-bit words of user-program space, on 64 tracks. (A maximum of 256 tracks were in theory available on each drum). Input and output was via 5-bit teleprinter equipment.

The digit frequency is 100 KHz, giving a 10 microsecond digit-period. The drum is synchronized to the processor's clock, allowing more than one drum to be added if required. 24 digit-periods (240 microseconds) are known as a *beat*. The basic machine rhythm consists of alternate *scan* and *action* beats. These are approximately equivalent to the modern *fetch* and *execute* phases except that a normal Ferranti Mark I and Mark I* instruction takes two *scans* and two *actions*, with an additional *action* beat if a 40-bit or 80-bit (rather than a 20-bit) arithmetic operation is called for. During the *scan* beats, the opportunity was taken to regenerate a line from the volatile CRT electrostatic store.

A four-beat instruction takes 960 microseconds, whilst a five-beat instruction takes 1,200 microseconds (ie 1.2 milliseconds). The timing of the principal computational instructions is as follows:

| | |
|--|---------------------------------|
| Simple arithmetical and logical functions: | 5 beats (total 1.2 millisecs.) |
| Multiplication instructions: | 9 beats (total 2.16 millisecs.) |
| Most other instructions: | 4 beats (total 0.96 millisec.) |

Both the Ferranti Mark I and the later Ferranti Mark I* computers are serial machines employing EF50 pentodes as their principal vacuum tubes (thermionic valves). Each computer contains approximately 1,600 pentodes and 2,000 thermionic diodes. The main CPU is contained in two bays, each 17 ft long by 9 ft high (5 metres x 2.7 metres) and consumed 25kW of power. At Manchester University, the computing room was electrically screened by a wire network that was built into the walls and continued over the windows and lights. This precaution was taken because of concerns about the sensitivity of the computer's electrostatic storage system to extraneous electro-magnetic noise such as was caused by the nearby City of Manchester trams!

The normal input/output equipment for both the Ferranti Mark I and Mark I* computers consisted of 5-track paper tape readers operating at 200 characters/second, paper tape punches operating at 15 characters/second, and a teleprinter printing at six characters/second. Some of the later machines also had a fast Bull lineprinter attached.

For both the Ferranti Mark I and the Ferranti Mark I*, the word length from a programmer's viewpoint is 40 bits. Instructions are 20 bits, packed two to a word. The smallest unit of addressable storage is a 'line' of 20 bits. Whether representing data, instructions or addresses, Ferranti adopted the convention that the most-significant bit is always held at the **right-most** end of the bit-pattern – ie 'backwards binary' – for the Mark I. This notation is explained more fully in Section F1X3. For the Mark I*, normal binary was used.

4. Summary of the main differences between the Ferranti Mark I and the Mark I*.

After delivering a second Ferranti Mark I computer to the University of Toronto in 1952, the Ferranti company made certain improvements to the basic design in the light of user experience. With the financial backing of the National Research Development Corporation, Ferranti produced the Mark I* (pronounced *Mark One Star*). This computer differed from its predecessor in several minor respects and in three major ones:

- (a) simplified instruction set, containing just 30 functions instead of 48 – (see section F1X3).
- (b) an extra 64 x 20 bits of fast CRT cache, used for buffering a fast lineprinter;
- (c) numbers in the Ferranti Mark I* were written down and displayed in conventional binary, rather than the ‘backwards binary’ of the Mark I.

Physically, the two types of computer look very similar. There are two main clues to distinguishing the types. Firstly, the operator’s console for the Ferranti Mark I* is simpler in layout; it has just one large and one small display tube, switchable so as to display the contents of a chosen tube of primary memory and one of {A, B, C, D}. In contrast, the Mark I’s console has two large and four small display tubes. Secondly, each logic door of the Mark I contained six smaller heater transformers, whereas a Mark I* logic door contained three larger transformers in square screening boxes.

As far as input/output equipment is concerned, a later addition was the attachment of a Bull high-speed printer (or *lineprinter*). This prints up to 64 characters per line on a roll of paper, the paper being automatically advanced by one line when each line (of up to 64 characters) has been printed. On the Ferranti Mark I*, the characters to be printed are first loaded into the 64 lines of the computer’s extra *output tube* – which is a CRT memory identical in form to each of the tubes in the main primary memory.

In operation, the differences between the two types of machine are best described with reference to the programmer’s viewpoint – see section F1X3.