

Appendix 3

Technical Details of the Elliott 400 Series Computers

A3.1 General Characteristics of the Family

Borehamwood produced four types of computers in the 400 series: the 401, 402, 403 and 405. A paper specification for an Elliott 404 was produced in 1956 [1] but the design never materialised. In any case, the architecture was very similar to that of the 405. The 404's proposed instruction format had one extra bit in the Op code field and one less in the address-field when compared with the 405's format. The instruction set of the 404 was similar in scope to that of the 405 – although the 404 would have allowed for the double modification of operand addresses.

The 400-series machines all had 2's complement serial arithmetic units employing thermionic valves (mostly double-triodes). All machines in the range used basically the same family of Elliott packaged circuits – with some additions to the family as time went on. The rhythm of the CPU was clocked to a timing track on the drum or disc backing store. Central registers and fast storage (where provided) were implemented as nickel magnetostrictive delay lines. Backing storage was via fixed-head magnetic drums or discs.

The Elliott 401 went through a number of modifications in its lifetime, of which the few details to have come to light are presented later, in Sect. A3.3. The description given below is for the original (1953) version of the machine. The description of the 403 has not been easy to verify but is believed to be correct. The descriptions of the 402 (also known as the 402E) and of the 405 are believed to be typical of standard production versions.

Only two machines of the Elliott 400 series had hardware floating-point facilities. These were known as the 402F variant, first delivered to Germany in 1958, for which the floating-point arithmetic unit added a further three cabinets to the normal complement of six cabinets. (Laurence Clarke, who designed the floating-point unit, remembers the floating-point unit as only one extra cabinet [2], but it depends upon whether one only counts fully populated logic cabinets or includes cabinets containing power supplies, cooling fans, etc. See, for example, the photograph in Fig. 5.6). One of the Elliott 402F computers was put to good use for 12 years for optical lens design by the Leitz company, before being retired in 1970 [3]. A second

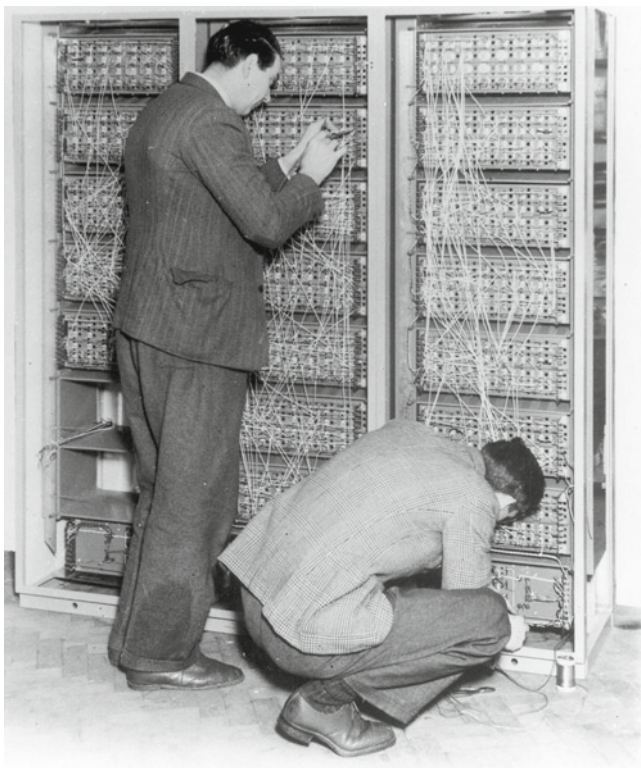


Fig. A3.1 The photo shows Computing Division staff in 1952, working on the backwiring of the three main racks which formed the CPU of the Elliott 401 computer. It may be deduced that, for Borehamwood, the 401 was more of a prototype than a production machine

402F went to British Rail. Delivery lists for all Elliott 400 series computers are given in Appendix 8.

The general properties of the 400 series computers are compared in Table A3.1. This table should be read in conjunction with the further details that follow. Note, however, that the surviving original Elliott source documentation is not always consistent, specifications evolving during the production life of each type of machine. Note especially that the fine details of the Elliott 401 and its peripheral equipment underwent several improvements during the years 1953–1965, first at Cambridge and then at Rothamsted – see Sect. A3.3 below and Chap. 5. A good first-hand account of the engineering environment in which the 400 series machines were developed at Borehamwood is given by Laurence Clarke in [4].

All the 400 series machines shared the same logic circuits, based on an original design by Charles Owen but converted by Laurence Clarke in 1952 to use the recently-introduced 12AT7 triodes instead of the more expensive (but more reliable) miniature pentodes. Laurence has described the circuit design philosophy as follows [4]. ‘The Owen circuits were based on a delay circuit which was very

Table A3.1 General characteristics of Elliott 400 series computers. The figures should be interpreted in the light of the comments given in the text

	401	402E and 402 F	403 (WREDAC)	405
Word length, visible bits	32	32	34	32
Word length incl. gap bits, if any	34	34	34	34
Digit period, microseconds	3	3	3	3
Instruction length	32	32	17	16
Instruction format, addresses	1 + 1	1 + 1	1	1
No. of instructions/word	1	1	2	2
Visible central registers (apart from Acc, Acc extn., Multiplier and handkeys)	2 (X, Y)	7 B-lines	3 B-lines chosen from 12 fast registers	3 B-lines
Fixed-point ADD time, min., milliseconds	0.204	0.204	0.204	0.306
Fixed-point MPY time, min., milliseconds	7	3.3	3.3	3.3
Max. Primary store (<i>Immediate Access</i> plus <i>Quick Access</i>), words	(none)	15 words	512 words	512 words
Backing store, words, max.	2,944	3K	16K	32K
Typical power requirements	5 kVA	7 kVA (11 kVA for 402F)	15 kVA, plus ?? kVA for WREDOC	?
Date first delivered	1953, 1954	1955	1955	1956

determined. The input was clocked; this caused a valve to make an inductance, which then charged up a condenser. At the end of the next clock period a reset pulse sucked the charge out very rapidly, and that left you with a totally determined pulse. What this enabled us to do was to put it together in very standard forms [ie packages] with a number of other devices – AND gate, OR-gate, digit delay, coincidence gated converter. These were combined to produce a gated delay – the standard circuit which you could then turn into what in other computers of the day was a flip-flop.

‘I think this activity was the beginning of the end of suck-it-and see electronics ... At Elliotts we started to work in a more methodical fashion. We analysed worst-worst cases for the logic loading rules ...’. This led to standard logic packages of predictable performance which could be combined without the logic designer having to worry about the internal circuit details.

A3.2 Elliott 401: Details of Fast Storage and the Instruction Set

The definitive specification of the 401 as it emerged from Borehamwood in 1953 is given in [5]. Note that a set of original 401 drawings was deposited in the Science Museum in London in 1991 [6]. The computer itself, as it existed in its last working state in 1965, is physically preserved in a Science Museum store at the time of writing. An original register-to-register diagram of the 401, and the details of the address-generation mechanism using modern terminology, are given in Figs. A3.2 and A3.3.

The Elliott 401’s main (and only) store was a disc (strictly speaking, a fixed-head *drum*, in more recent terminology). When track switching occurred, and when information was written to the disc, the reading amplifiers were temporarily ‘paralysed’ for three or four word-times. If a program requested a read during a paralysis period, hardware automatically delayed reading for a complete revolution of the disk. One revolution-time equalled 13.1 ms. The 401’s CPU had a two-beat fetch-execute rhythm. That is to say, instructions were selected and obeyed alternatively. The basic *beat* was 34 digit-times of 3 μ s each (thus allowing for the two-digit gap). The total fetch-execute sequence was therefore basically 204 μ s and the basic fixed-point add time was also 204 μ s. However, two-successive *ADD* instructions could only be executed at this rate if optimum programming had been used to position these two instructions appropriately on the disk. If the second instruction was ‘missed’, the disk may have to complete a revolution (13.1 ms) before the required instruction became available. In view of this, and disk paralysis period noted above, it seems reasonable to take an average add time for practical programs of perhaps 3 ms if attempting to compare the 401 with the instruction speeds of other contemporary machines with random-access primary memory. This touches on the art of *optimum programming*, for which the precise location of instructions and operands in a sequentially-accessed memory was important.

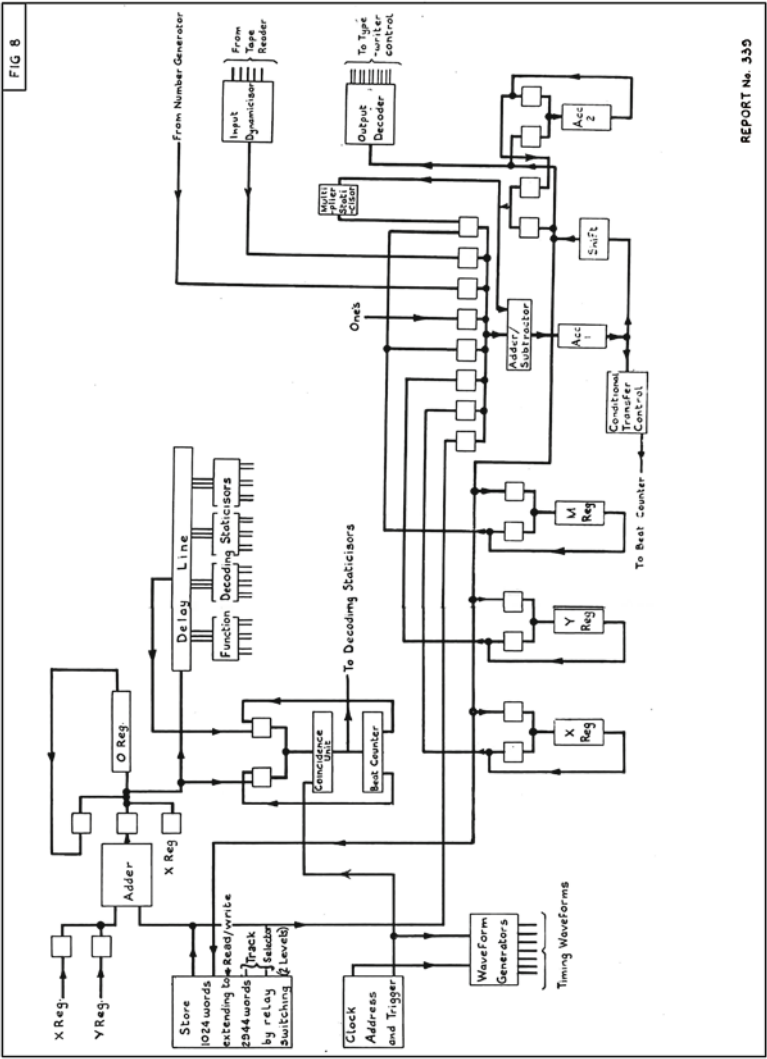


Fig. A3.2 An original diagram, taken from a 1954 Borehamwood research report, shows the layout of the arithmetic unit and central registers of the Elliott 401 computer. The main drum memory is to the left of the diagram and the input/output equipment is on the upper right. The adder/subtractor unit is on the lower right

REPORT No. 339

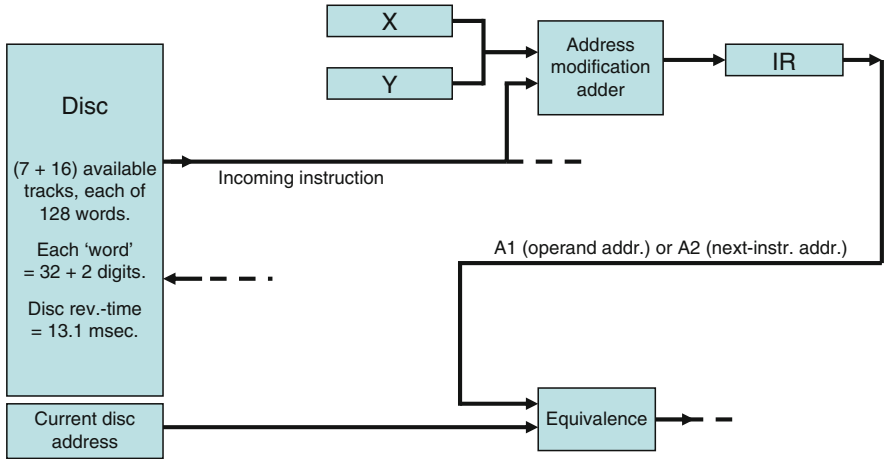


Fig. A3.3 This is a simplified representation of the 401’s address-generation mechanism, based on the details shown in the upper-right part of the original diagram in Fig. A3.2. Because the main disc memory was sequential in nature, information could only be read or written when equivalence occurred between the disc’s current address and the desired operand-address or instruction-address

The general layout for the Elliott 401’s 32-bit instruction was as follows:

10	3	3	3	3	10
A2	S	F	D	C	A1
Address of next instruction	Acc source	Op code	Acc dest.	Control (K)	Operand address (also called <i>timing</i>)

The 10-bit addresses A1 and A2 refer to locations on the 401’s disc store. A 10-bit address is formed as three track-digits followed by seven location-digits. The {S, F, D, C} bits give about 28 useful operations via the following combinations:

	S	F	D	C
000	Nowhere	ADD	Nowhere	Normal
001	All ones	MPY	Round off	Test acc negative
010	Input	Left shift n	Output	Test acc zero
011	R3	Shift right n	R3	Use R3 to modify
100	R4	LOAD	R4	Use R4 to modify
101	R5	NEQ	R5	Use R5 to modify
110	Memory	SUB	Memory	Double length
111	Handkeys	AND	Track 7 switch	Optional stop

Some combinations are forbidden, for example {6, –, 6, –} which attempts to specify memory as both the source and destination. The coding of the S, F, D and C fields as given above differs in minor respects from that shown in the original Borehamwood report [5]. The names of the central registers also evolved over the

years. Some equivalent names are as follows: R1 = accumulator; R2 = accumulator extension; R3 = Multiplier register, M; R4 = general register X; R5 = general register Y. The main computer, including power supplies and cooling fans, occupied seven cabinets.

Further explanation of some of the Elliott 401 instructions is now given. The *NEQ* instruction is the logical *Not Equivalence* operation, also called *Exclusive OR*. The Handkeys signify a 32-bit value set up manually via switches, originally called the *number generator* by Borehamwood. When shifting the accumulator up or down arithmetically by n places, the value of n is given by the difference ($A2 - A1$). During multiplication, the contents of R2 are multiplied by the contents of R3, leaving the product (62 digits plus a sign digit) in R1 (more significant half) and R2 (less-significant half). If only the most significant half of a product is required, round off may be used. The product, rounded to 31 binary places, is then placed in R1. The time for multiplication given in Table A3.1 is that for the post-1954 version of the 401 – (the original time was quoted as being ‘between 7 and 10 milliseconds, depending upon sign’). When using the *Double Length* indicator, R1 and R2 together form a two-word accumulator. When using the *Optional Stop* indicator, the computer will halt before obeying the instruction and wait until the operator presses the single-shot key on the control console. The machine can be set to ignore *Optional Stop* instructions by depressing a switch on the console.

The original disc store for the Elliott 401 ‘comprised 8 tracks closely spaced near the outer rim of a 9-inch diameter disc, each track carrying 128 words’. This was soon increased to 23 tracks, each holding 128 words. Since only ten address bits are provided in the 401’s instruction, only eight tracks are available to the programmer at any one time. Tracks 0–6 are by convention fixed tracks, and ‘track 7’ can dynamically be made to refer to any of the 16 remaining tracks, logically numbered 0–15, by means of the *Track 7* instruction. When switching between one *track 7* and another *track 7*, the programmer needed to allow for a full drum revolution of approximately 13 ms.

Selecting the *Input* combination of the S-bits in an Elliott 401 instruction causes a row to be read from the 5-track Tape Reader and the tape to be advanced by one row. The integer represented by the 5-bit combination is placed in the least-significant end of the accumulator. When the *Output* combination of the D-bits is selected, either a paper tape punch or a typewriter (teleprinter) is activated, the choice being determined by the setting of a manual switch on the equipment. In either case, the least-significant five bits of the accumulator are either punched on 5-track paper tape or printed as an alpha-numeric character on a page.

A3.3 The 401 at Rothamsted: Programming and Enhancements

In Tables A3.2 and A3.3 is an example of a library subroutine, as used on the Elliott 401 computer at Rothamsted. The descriptions come from a typewritten computer manual at Rothamsted [7]. The author/date of this particular subroutine is given as M J R Healy, 7 October 1954 and the code implements the square root function.

Table A3.2 The formal description of the square root library subroutine

Sub-routine
Square root
 $y = \sqrt{x}$, $0 \leq x < 1$
Addresses occupied: V.64 - V.82
Registers used: All.
Enter at V64, link in R1, x in R4
Result: y in R1 (0 in R4)
Time: 0.1 - 1 sec. approx, longest for small numbers.
Accuracy: The result has two fewer significant (binary) figures than the original number.
Method: see EDSAC sub-routine S2
 $a_{n+1} = a_n (1 - \frac{1}{2} c_n)$ $a_0 = x$ $a_n \rightarrow \sqrt{x}$
 $c_{n+1} = \frac{1}{4} c_n^2 (c_n - 3)$ $c_0 = x - 1$ $c_n \rightarrow 0$.
 M.J.R.Healy
 7 October 1954

Table A3.3 The code for the square root library subroutine

Programme
 Square root
 V.64 V.82 4460 V.65
 V.82 V.81 0002 V.71 If x = 0, jump to end of programme.
 V.81 V.72 6000 V.70 c_0
 | \rightarrow V.72 V.75 0330 V.73 $\frac{1}{4}c_n$; c_n to R3
 | V.75 V.77 6600 V.76 $\frac{1}{4}c_n - \frac{3}{4}$
 | V.77 V.79 4456 V.78
 | V.79 V.66 0110 V.35 $a_n c_n$
 | V.66 V.68 0300 V.67 $\frac{1}{2}a_n c_n$
 | V.68 V.73 4440 V.69
 | V.73 V.78 4600 V.74 a_{n+1}
 | V.78 V.80 5446 V.79
 | V.80 V.67 0110 V.36
 | V.67 V.74 0006 V.68
 | V.74 V.69 0110 V.38 c_{n+1}
 | V.69 V.72 0002 V.71
 V.71 V.65 4440 V.72
 V.70 4.00 0000 0.00 || -1
 V.76 3.00 0000 0.00 || $\frac{3}{4}$
 V.65 0.12 3456 7.89 || link space.

The notation in Table A3.3 is as follows, where <SFDC> are the four octal digits giving the {acc. source, op code, acc. destination and control}:

<addr. of instr.> <addr.of next instr.> <SFDC> <addr. of operand>

Addresses are given as decimal numbers preceded by V. Table A3.2 contains the original description as found in the Rothamsted manual and Table A3.3 the corresponding program code.

A revealing snap-shot of life as an early user of an early computer is given by Gavin Ross, who programmed the Elliott 401 in machine code at Rothamsted [8]. Gavin writes as follows: ‘Frank Yates [Head of Statistics] obtained the agreement that he could use the Elliott 401 for whatever purposes he and his staff thought fit, and although the first priority was to take over the routine analysis of field experiments and surveys, the statisticians who learned to program it wrote more general routines for linear regression, multivariate analysis, nonlinear modelling, biological assay and cluster analysis. Elementary functions like division, square roots, logs and exponentials had to be programmed as subroutines, later supplemented by statistical functions like the Normal Probability Integral. We had to devise our own routines for sorting lists of numbers, for finding our way round multiway tables, for replacing a rectangle of rows and columns by a rectangle of columns and rows, for cleaning up the store, and for numerical analysis procedures such as inverting matrices and finding eigenvalues and eigenvectors, performing integrals and differentials and optimising functions of several variables. These routines died with the machine, and we had to start again with the next computer [a Ferranti Orion].’

The Elliott 401 remained at the Rothamsted Research establishment from 1954 until 1965. As far as can be ascertained [7], the main changes to the hardware during that time were as follows:

- Replacing the character code for digits by a parity-checked code, in 1955.
- A high-speed paper tape reader (a Ferranti Mark II reader) was installed in 1956.
- Replacing the typewriter by a tape punch, to be read by a teleprinter off-line.
- Installing a card reader which was linked to the 401’s hand-switches, in 1956. This read 32 columns from 80, at a rate of about 15 cards per minute.
- Adding a controlled sequential addition facility in 1958. This formed the sum of a specified number of consecutive numbers stored on a given track, starting at A1 and ending at A2. Gavin Ross remembers [8] that ‘It was supposed to speed things up, but I do not recall it being used by programmers.’
- Adding an extra set of immediate access registers S3–S5, usable instead of R3–R5 without losing information, in 1958.
- The disk was replaced by a drum in 1960.
- Track 6 switching, thereby allowing eight more tracks to become available, was implemented in 1961.

A3.4 Elliott 402: Fast Storage and Instruction Set Details

The Elliott 402 had 15 *Immediate-Access* registers, implemented as single-word nickel delay lines, and faster (electronic) switching of disc tracks. Therefore, the 402’s *average* add time would have been faster than that of the Elliott 401. *Immediate-Access* locations 0–7 could be used as index registers (B-lines), with the convention that location 0 always held the value zero.

The Elliott 402’s 32-bit instruction had a similar format to that of the 401 given earlier. As before, the format included ten bits of operand-address and ten bits for specifying the address of the next instruction. However, with the 402 the first 16 addresses 1–15 were mapped onto 15 fast registers called the *Immediate-Access Store* and the handkeys were mapped onto address 0. This eliminated the need to refer explicitly to the 401’s named registers R4 (X) and R5 (Y) and thus allowed the S-field and the D-field to be reduced in size. The remaining 12 bits of the 402’s instruction were re-allocated into five fields, of which the B-field allowed seven of the 15 fast registers to be used for address modification (B = 0 indicating ‘no modification’). The layout was as follows.

10	3	3	2	2	10	3
A2	S	F	D	C	A1	B
Next-instruction address	Acc source	Op code	Acc dest	Control	Operand address	Specifies which of the 7 IAS locations is to be used for modification

Actually, the fields were physically arranged in a slightly different order as follows, assuming the least-significant end of the word is at the right-hand side:

<A2> <S> <F> <D> <A1> <C>

The new repertoire of ALU operations was defined for the Elliott 402 as follows:

	S	F	D	C
000	Zero	ADD	Nowhere	Normal
001	M register	MPY or DIV	M register	Test negative or Divide
010	Memory	Left shift n	Memory	Count
011	Input	Shift right n	Output	Track 7 switch
100		LOAD		
101		AND		
110		SUB		
111		Negate		

As compared to the Elliott 401, the 402 had an improved multiplier design, multiplication taking 3 ms regardless of sign. It also had hardware division, also taking 3 ms, and an improved disc store.

A3.5 Elliott 403 (WREDAC): Fast Storage and Instruction Set Details

The Elliott 403 had a fast, four-word, instruction buffer so that decoding of a following instruction took place whilst the current instruction was being executed. Then the 403 had a comparatively large 512-word *Immediate-Access* store implemented as 12 single-word delay lines (effectively random-access) in the first

sub-section, followed by 127 delay lines each of four words. The average add-time was therefore considerably faster than either the Elliott 402 or 405. The Elliott 403 had three visible index registers (B-lines), pre-selected by program from four sets held in the 12 fastest storage locations. For B1 and B2, digits 20–32 are added to the current instruction, thus modifying the Function (op code) and operand-address fields. If B3 is specified, the address of the next instruction (not to be confused with the next word) is added to the current order's address digits.

If the least-significant digit, C, in an Elliott 403 instruction is zero then the instruction refers to arithmetic operations. When $C = 1$ *input/output transfer* instructions are specified. The format for arithmetic instructions ($C = 0$) is as follows:

5	9	2	1
F	A	B	C
Op code	Operand-address	Modifier register	Code = 0

There are 32 ALU instructions as specified by the F bits, most of them acting variously on the double-length and single-length accumulators. The accumulator-based repertoire includes:

Add; Subtract; Clear; Clear and Add; Clear and Subtract; Store;

Swap – (i.e. swap the contents of the accumulator with the contents of a store address);

AND; Logical shift left; Logical shift right; Multiply; Divide; Normalize.

Unlike the other computers in the Elliott 400 series, the 403 has three instructions relevant to multiplication, whose action is as follows:

SET MULTIPLIER REGISTER. This loads the multiplier register R with the contents of a specified memory location. This is the only means by which a programmer can refer to the multiplier register.

MULTIPLY and ADD. The contents of a specified memory location is multiplied by the number in the multiplier register R and the 66-digit double-length product is added to the accumulator.

MULTIPLY and SUBTRACT. As above, except that the product is subtracted from the accumulator.

Additionally with $C = 0$, the Elliott 403 has six *test and jump* instructions involving tests on the contents of the accumulator and two *test, jump and count* instructions involving the contents of modifier registers (B-lines). There are two instructions for loading a value from memory into a selected B-line. Each B-line was 17 bits (defined as the *even* half of a word).

Finally with $C = 0$, the Elliott 403 has a special *Use Logic* instruction. To quote [9]: ‘One special order is reserved in the 0-code orders to determine the mode of operation of the machine. Allowance has been made for up to 8 different modes of operation, each one of which may be sub-classified in 64 ways. Only one

mode has been attached to [i.e. implemented in] the machine to date, namely, the “use B-lines” mode’. This instruction configures various B-line options, namely:

- a. Specifies the group of three B-lines, from the four available groups, which will be switched to high speed store addresses 1,2,3 respectively;
- b. Specifies the B-line to which various *test-and-jump* instructions apply;
- c. Specifies the two B-lines which may be added together before being used as modifiers.

Looking back, it is not entirely clear what was in the minds of the Long Range Weapons Establishment and Borehamwood when this instruction was first defined. The option (a) is historically interesting since it hints at what was later to be known as *context-switching*. Option (c) gives double modification.

When $C = 1$, the machine obeys one of 32 Input/Output instructions as specified by the five F bits. The repertoire includes:

Four instructions for handling slow, 5-bit, input/output;

16 Instructions for handling magnetic tape input/output;

Two instructions for transferring 64-word blocks of data to/from the disc.

The reason why there were as many as 16 magnetic tape instructions is that, originally, the Elliott 403 had a single magnetic tape control channel. A second, independent, channel was then introduced. It was decided to dedicate new instructions to the second channel so as to preserve the integrity of pre-existing program code.

A3.6 Elliott 405: Details of Fast Storage and the Instruction Set

Figure 3.4 shows the overall architecture of an Elliott 405 system, indicating the main data paths. The Elliott 405’s CPU had a three-beat rhythm: fetch instruction, read/write from/to store, perform operation. Therefore, the minimum instruction time was 306 μ s. The Primary Store had two sections: a fast *Immediate-Access* section and a not-so-fast *Quick-Access* section. The *Immediate-Access Store (IAS)* consisted of either 4 or 20 one-word nickel delay lines. If there were 20 words of IAS, then the *Quick-Access* store had one of its 16-word delay lines replaced by 16 of these 1-word delay lines. The *Quick-Access Store* consisted of at least 20 16-word nickel delay lines (each of total circulation-time 1.6 ms) and up to a maximum of 32 such lines, giving a maximum of 512 words. The actual amount for any specific machine depended upon the physical configuration chosen at installation-time – see Chap. 10. Since it was quite possible to hold a small routine and its working-space in the immediate-access store, it might be thought reasonable to take 306 μ s as the Elliott 405’s average fixed-point add time. However, the story is not so simple, as is now explained.

The physical addressing of each 16-word delay line in the Elliott 405 is so arranged that the timings of successive words in the *Quick-Access* store differ by

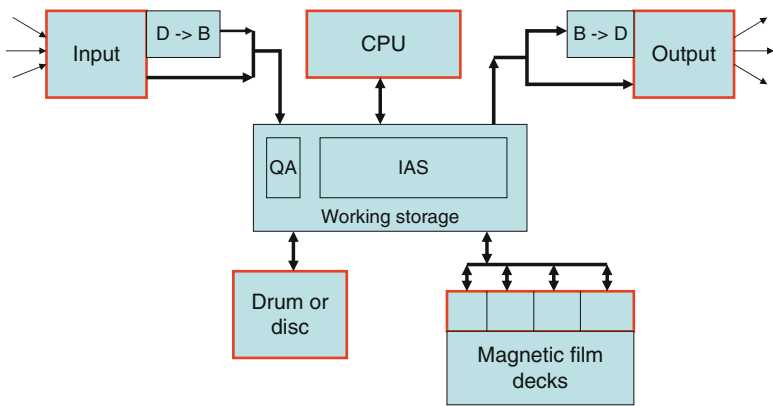


Fig. A3.4 Schematic diagram showing the main sub-sections of an Elliott 405 computer. The programmer could initiate autonomous activity on up to five logically separate local machines: one for input transfers, one for output transfers, one for drum or disc transfers, one for magnetic film transfers, and one for central processing

three word-times in the 16-word cycle. The total time taken by two successive instructions depends very much on the relative positions of the two instructions and their operands in these 16-word delay lines. For example, assuming that a pair of simple ADD instructions is stored at address 4, [10] gives two extreme cases to illustrate how the choice of operand-address can dramatically alter the total time taken by the pair:

	First instruction	Second instruction	Total time taken
Case (a)	Acc = Acc + addr(10)	Acc = Acc + addr(15)	3 word-times (total 0.918 ms)
Case (b)	Acc = Acc + addr(20)	Acc = Acc + addr(5)	35 word-times (total 10.71 ms)

Instructions for the Elliott 405 are packed two to a word. The machine obeys first the one in the most significant half (i.e. the one which occupies digit positions 17–32), and secondly the other which occupies positions 1–16. The Elliott 405 instructions were divided into two groups, distinguished by the setting of the least-significant bit, denoted as C. When C = 0 the format was:

4	9	2	1
F	A	B	C
Op code	Operand address	Modifier register	Code = 0

The Elliott 405 had 16 computational instructions, denoted as *0-codes*, for which C = 0. Briefly, these acted as follows:

- a. Operations between the contents of the accumulator and the contents of a memory address: Add, Subtract, Reverse subtract, Multiply, Divide, AND, Swap (i.e. exchange), Load. Note that both single- and double length multiplication

are available. For double length, the remaining 31 digits of the full product are left in *Immediate-Access Store* location 3. The type of multiplication is set by a *I-Code* instruction – see below. Once double length working has been set, all subsequent multiplications will be double length until another order resets to single length working. Multiplication instructions take 33 word-times. The Elliott manuals contained a special warning for the Divide instruction. Quoting [10], ‘the logical process involved requires that the divisor be or be made greater in absolute magnitude than the dividend before division takes place. The result may be in error by up to 2^{-31} so that if the actual divisor is a factor of the dividend, the result may be slightly different from the true result. This order takes 33 word times.’

- b. Operations on the accumulator involving a constant (literal): Logical left shift n times ($1 \leq n \leq 16$); Logical right shift n times ($1 \geq n \leq 16$); Multiply by n (where $n = 0, 2, 4, 8, 10, 12$). This last instruction is designed to speed number-conversion in decimal or sterling calculations.
- c. Control transfer instructions: unconditional absolute jump; absolute jump if Accumulator is negative; absolute jump if the contents of address (1) is zero and add 2^{-12} to address (1) in any case.
- d. Move instructions: these permit any block of 16 words to be moved to any other block-position in primary memory, the size/position being chosen so as to correspond to one of the 16-word nickel delay lines making up the memory. One *Move* instruction covered the address-range 0–255, another *Move* instruction covering the range 256–511. There was an oddity if the lowest block (‘line 0’) was specified. As explained in [10], locations 0–3 of block 0 are not part of the *Quick-Access Store*, since addresses 0–3 designate the number generator (0) and *Immediate-Access Store* locations 1–3. These four addresses are known as ‘ghost locations’. (If the contents of any other block are transferred to block 0, the first four words pass into these ghost locations and remain there unaltered until the whole block is again transferred elsewhere. The only arithmetic functions in which the contents of the ghost locations can take part immediately are Multiplication and Division).

When $C = 1$, the so-called *I-Code* instructions mostly concerned input/output transfers and backing-store activity. An exception is a group of three *I-Code* instructions that are strictly part of normal computational activity. These three are placed in a group known as *Overall Machine Control* in the Elliott 405 literature, for which the S-bits (see below) are set to the value 7. The three instructions are:

Set single length multiplication.

Set double length multiplication.

Set link. This stored the contents of the Program Counter (called the *Sequence Control Register*) in memory location 1.

The general format for *I-Code* instructions for the Elliott 405 is as follows:

6	3	3	3	1
N	T	F	S	C
Code = 1				

The S-bits are assigned as follows:

S = 0: a group of instructions for handling input from paper tape and punched cards.

S = 1: a group of instructions for character-based output to paper tape, typewriter or magnetic film.

S = 2: a group of instructions for block transfers (64 words) to/from disc or drum.

S = 3: a group of instructions for block transfers to/from magnetic film.

S = 7: the three *overall machine control* instructions described above.

As can be seen, only some of the possible <N, T, F, S, C> combinations are assigned, thus allowing for future expansion if new devices were to be introduced for the Elliott 405.

A comparison of the relative effectiveness of the Elliott 405's instruction set, when compared to that of its market rival the IBM 705, is given in Sect. A3.9 when discussing a large payroll application for the British army.

A3.7 Elliott 405: General Configurations and Disc Options

The Elliott 400 series computers used fixed-head magnetic drums though, because of their physical shape, they were often called discs. The physical characteristics given in Table A3.4 are for the Elliott 405 computer. However, the 'drum' column is similar to the backing-store arrangement for the Elliott 402 and the 'small disc' column is similar to that of the Elliott 403.

Physically, an Elliott 405 consisted of a number of cabinets [10, 11], each measuring about 6 ft 6 in. high on a 2 ft square base (about 2 m high by 0.6 m²) and each weighing about 320 lb (145 kg). Apart from the operator's console, an installation could be build up from the following units, where QAS signifies Quick-Access Storage composed of 16-word nickel delay lines, of which there could be a maximum of 512 words spread throughout the whole computer – (this was in addition to either 4 or 20 single-word lines of Immediate-Access Storage (RAM)).

Table A3.4 Backing-store options for the Elliott 405 computer. A larger disc of capacity 32K words was also available

	Drum	Small disc
Diameter (in.)	8.5	19.25
Width (in.)	1.5	0.5
RPM	4,600	2,300
Tracks	32	64
Sectors	64	256
Words	4,096	16,384
1 rev. time (ms)	13.2	26.4
Packing density (digits/in.)	166	166

Unit and description	No. of cabinets
System Centre (CPU and 128 words of IAS)	2
Drum Store, 4K words (including 128 words of IAS buffer)	2
Disc Store, either 16K or 32K words (incl. 128 words of IAS)	3
35-mm Film Store – Master (incl. 64 words of IAS and 2 decks)	3
35-mm Film Store – Slave (including one film deck)	1
Simple Input/Output (character-at-a-time)	1
Input/Output Compiler (including 16+ words of IAS)	2
Power Supply – master unit	1
Power Supply – slave unit	1

In [12] a small 405 system, such as the one first delivered to Norwich City Council [13], was stated to have nine cabinets and a control console. A medium installation, such as the one used in 1957 at Borehamwood for in-house applications, is believed to have had at least 15 cabinets plus console. In theory the largest Elliott 405 could have 82 cabinets, most of which would be magnetic film units.

An Elliott 405 could have up to four magnetic film *Master* units, controlling a total of up to 16 decks. In addition, there was an off-line output device (called MUF-PUP or MUFPT) which allowed a magnetic film deck to write at high speed to either a lineprinter or to a cluster of paper tape punches or to punched cards. The more demanding, high-performance, lineprinters available for the 405 printed 300 lines/min, with 140 characters per line.

In conclusion, a complete 405 installation gave the programmer several autonomous, programmable, local machines – one for central processing, one for disc or drum memory, one for magnetic film, one for bulk input/output – though it was rare for all to be actively engaged concurrently. The whole assembly of facilities allowed the movement of large amounts of commercial-type data to be overlapped with computation.

A3.8 Input/Output and Bulk Storage for the Elliott 400 Series

The properties of the common input/output devices for Elliott 400 series computers evolved to some extent over the relevant period 1953–1962. The original input/output medium for all 400 series computers was 5-track paper tape, based historically on Creed teleprinter equipment. The five bits could sometimes be treated by a computer as four information bits and a parity bit – for example in the encoding of numerals. Different manufacturers used different 5-track conventions. The complete Elliott coding of alphanumeric characters, called Elliott Telecode, is given in Table A3.5 for the Elliott 402, 403 and 405 computers. It may be seen that the numerals all have odd parity for the 402 and 405 conventions. Aside from the desirability of maintaining a conveniently ordered collating sequence for letters and numbers, Table A3.5 indicates that there is some choice in assigning visible characters to bit-patterns.

Table A3.5 Teleprinter codes for various Elliott 400 series computers. The 1's in a telecode character indicate holes in paper tape and the 0's no holes

Telecode character	Letter shift	Figure shift, 402	Figure shift, 403	Figure shift, 405
00000	Blank	Blank	Blank	Blank
00001	A	1	.	1
00010	B	2	*	2
00011	C	*	1	*
00100	D	4	=	4
00101	E	\$	2	\$ or “
00110	F	=	3	= or £
00111	G	7	;	7
01000	H	8	.	8
01001	I	‘	4	‘
01010	J	,	5	,
01011	K	+	:	+ or 11
01100	L	:	6	:
01101	M	-	%	-
01110	N	.	;	.
01111	O	%	(%
10000	P	0	-	0
10001	Q	(7	(
10010	R)	8)
10011	S	3	?	3
10100	T	?	9	?
10101	U	5	/	5
10110	V	6	+	6
10111	W	/)	/
11000	X	@	0	@ or &
11001	Y	9	£	9
11010	Z	£	@	£ or 10
11011	Figure shift	Figure shift	Figure shift	Figure shift
11100	Space	Space	Space	Space
11101	Carriage return	Carriage return	Carriage return	Carriage return
11110	Line feed	Line feed	Line feed	Line feed
11111	Letter shift	Letter shift	Letter shift	Letter shift

Punched card equipment for input/output was subsequently introduced for the 402 and 405 computers. The 401 was also later equipped with a reduced form of punched card input via an adaptation of the 401's *handkeys* combination of the S-bits of an instruction. For this, any 32 columns out of 80, as selected by means of a plug-board, could be transferred to the accumulator at the rate of 100 cards per minute.

For 400 series computers, the size of a *character* was defined according to the programming context, values of 4, 5 and 6 bits occurring at various points in the Elliott technical literature. There was no generally accepted unit such as an 8-bit *byte* in the 1950s. Practical data transfer rates are dependent upon particular devices when attached to particular computers and it is not easy to generalise and compare performances.

Early paper tape readers for the 400 series computers operated at about 40 chars/s (characters per second), to be followed by the Ferranti paper tape reader at 100 chars/s and, later still, by devices operating at 180 chars/s. In 1958 Borehamwood designed Elliott's fast paper tape reader, capable of transferring information at 1,000 chars/s. Card readers transferred data at the rate of up to 400 eighty-column cards per minute.

The original output teleprinters printed at about 10 chars/s. Paper tape punches worked at about 25 chars/s. Two forms of lineprinter were offered for the Elliott 405. Early machines typically had Bull lineprinters, giving 150 lines per minute with 92 characters per line. Later printers gave 300 lines per minute with 140 characters per line. As mentioned in Chap. 10, these lineprinters were usually driven off-line from magnetic film units because of their demands for high data-rates.

The Elliott 403 had quarter-inch magnetic tape decks and the 405 had 35-mm magnetic film decks, as described in detail in Chap. 10.

We end this account of Borehamwood's 400 series computers by describing the head-to-head evaluation of the Elliott 405 and the IBM 705 in 1956. In that year the Royal Army Pay Corps (RAPC) conducted a procurement exercise to select a digital computing system to replace the hundreds of electro-mechanical desktop machines that were used in the payroll calculations for hundreds of thousands of British soldiers. After sifting through the offerings of 14 UK and US computer manufacturers, the RAPC decided that the only two-candidate EDP systems capable of handling the huge payroll task were the IBM 705 and the Elliott 405. The background to this decision is outlined in Chap. 9. We now continue the detailed comparison of these two computers, based on the contemporary technical evidence contained in [14, 15].

A3.9 Comparison of the Elliott 405 and the IBM 705: The RAPC Procurement

A3.9.1 The Programming Task

The payroll tasks selected by the Royal Army Pay Corps (RAPC) Working Party for automation included [14, 15]:

- Maintenance of records of the soldiers' pay entitlements, tax deductions, dependants' allowances, automatic increments, etc.;
- Preparation of notifications to individuals and to military units of payment details and budget analysis;
- Keeping an up-to-date record of vouchers, army allowance order books, claims and correspondence;
- Production of public audit figures, management statistics and analyses, and internal system checks.

It was planned that input to the computer would be by punched cards, records would be maintained centrally on magnetic tape, and output would be in the form of printed notifications to the military units worldwide. The Working Party decided