Version 3, November 2011 **E3X2**

Systems architectures for the Elliott 800 series and 503 computers.

Note: all references are listed in section E3/X5.

General introduction to the Elliott 800 series and the 503.
The 800 series of computers were small or medium-sized, serial, transistorised machines that achieved a relatively high market penetration in the UK in the 1960s. The 802 was originally designed for two purposes: “to perform as a small general-purpose computer, and also to act as the computing centre for on-line process-control systems” (reference 4). The majority of installed computers were in fact the 803B variant. The complete genealogy is as follows:

801 An in-house test-rig; only one built (about 1957); used transistors-plus-cores circuits.

802 The production version of the 801. The 802 was launched in November 1958 at the Automation Exhibition at Olympia, London. It was a 33-bit word serial arithmetic machine, superseded about a year later by the 803A (see below). To minimise hardware the 802 held all its internal registers as a single bit-stream which circulated through the various stages of processing once per instruction. The 802 had a 'big L-shaped desk' format. The 802 used ferrite core logic. It was mostly transistorised, but valves were used for driving the core store and for triggering the core logic elements because transistors of the required power capability had not yet become generally available.

803A The fully-transistorised version of the 802, with the word-length increased to 39 bits. The first 803 worked in 1959. It had a cabinet-based design. The bit-stream architecture of the 803A was similar to that of the 802. The 803A was a short-lived design, superseded by the 803B. The 803A was in general rather slower than the 802, on account of the longer word-length. This was the spur to designing an improved version, the 803B.

803B The enhanced 803A, first working in 1960. The 803B employed more parallel paths (separate registers) internally instead of the bit-stream approach, and had hardware floating-point. The majority of '803s' sold would have been 803Bs. The selling price (1960 values) was £29,000. The 803B was said to have been "an immediate success. It was small and simple and completely transistorised. As a result, the power consumption of the basic machine was only 3.6KW. The transistor/core logic proved extremely reliable" (reference 8).

804 Paper design only; **805** Paper design only.

503. The larger and much faster parallel-arithmetic 503 was software-compatible with the 803. By the time Borehamwood came to design the Elliott 503, the characteristics of the available transistors had improved so the 503’s logic circuits were based on transistors (for

amplification) and diodes for the logic functions, rather that the earlier transistor-plus-core combination. The first production 503 was delivered in 1963 – (see section E4/X1). It was about 70 times faster than the Elliott 803.

The great commercial success of the Elliott 800 series during the early 1960s has been attributed to: (a) Low cost; (b) Ease of use and installation (no air conditioning required, powered from a standard electrical socket); (c) The core store and core logic proved to be extremely reliable.

The overall characteristics of the Elliott 800 series and 503 computers are compatred in the following two Tables.

	802	***803A***	***803B***	***503***
Word length, visible bits	33	39	39	39
Digit period, microseconds	6	6	6	?
Instruction length	16	19	19	19
Instruction format, addresses	1	1	1	1
No. of instructions/word	2	2	2	2
Basic ALU: serial or parallel?	serial	serial	serial	parallel
Fixed-point ADD time, unmodified, minimum, microseconds	612	720	576	7.2
Fixed-point ADD time, modified, max., microseconds	612	720	576	10.8
Fixed-point MPY time, min., milliseconds or microseconds	21.4 millisecs	29.5 millisecs	1.15 millisecs	34.8 microsecs
Hardware floating-point ADD time, min., microsecs.	no	?	864	13.2
Max. Primary store (core), words	1K	4K ?	8K	8K
Core backing store, words, max.	None	None	None	128K
Typical power requirements	2KVA	1 KVA	3.5 KVA	?
Date first delivered	1958	1960	1960	1963

General characteristics of the Elliott800/503 family of computers.

	Elliott 803	***Elliott 503***
Paper tape reader	500 ch/sec	1,000 ch/sec
Paper tape punch	100 ch/sec	100 ch/sec
Card reader	340 cards/min	340
Card punch	100 cards/min	100 cards/min
Lineprinter	300 lines/min	1,000 lines/min
Core backing store	none	71.9 micro/word
Magnetic film	4,350 ch/sec	Not applicable
Magnetic tape	Not applicable	42,000 ch/sec

Maximum transfer rates for a selection of input/output devices.

Elliott 802 systems architecture.

Underlying CPU technology. The 802 represented Elliott's first full-scale use of germanium junction transistors. In the mid-1950s, junction transistors offered the following advantages over thermionic vacuum tubes (valves): greater efficiency (eg lower power consumption); smaller physical size; greater reliability; cheaper cost. However, at the time when the 801 test rig was being built, the performance of available transistors fell below that of valves in respect of high-power and high-speed switching capabilities.

The 802's principle requirement for power circuitry concerned the main memory (ferrite core store) drive circuits and drivers for resetting the logic cores; for both these requirements, valves were accordingly used. Transistor characteristics were not sufficiently consistent for logic switching unaided, so Borehamwood used ferrite cores as logic elements and an OC71 transistor for amplification. These two components, plus a resistor, were mounted on 'daughter' printed-circuit boards measuring 1.5 inches x 2.25 inches (later reduced to 1.5 x 1.5 inches). About 30 daughter boards were then mounted upon Mother boards measuring 13 inches x 8 inches (about 33 cms x 20 cms). Suitable windings round each core were used to give two basic logic functions: (*A or* B) and (A *&* not B). Since a logic core had to be in a reset state at the start of each logical operation and then interrogated at the end of an operation, a two-part system was used whereby the state of one core was copied to a following 'slave' core. This yielded a basic digit-period of (3 + 3) = 6 microseconds. Logic could be performed on each of the two parts (*master* and *slave*) of a complete circuit module, resulting in some gains in economy.

Central registers and main memory. In order to keep the 802's hardware simple, a nickel delay line is employed as a single, 102-bit, *Operations Register* containing the following programmer-visible registers:

Accumulator: 33 bits;
Instruction-pair (Order Register): 33 bits;
Program counter (known as SCR): 11 bits.

The remaining 25 digit-positions in the Operations Register are used for hardware timing purposes. SCR contains the address of the current half-word instruction. Each time an instruction is obeyed, SCR is incremented. There is also a 65-digit double-length accumulator used in multiplication, division and shifting instruction. The least-significant 32 bits of the double-length register are called the A*uxiliary Register, AR*, and are visible to the programmer; the upper 33 bits of the double length register are inaccessible. Finally, there is an overflow register. One circulation of the 102-bit *Operations Register* defines the basic timing of the 802. Since the digit-period is 6 microseconds, the execution-time of simple instructions is 612 microseconds.

Two 16-bit single-address instructions occupy one 33-bit word – (see also section E3X3). They are separated by a one-bit marker called the *B-line*. If the B-line is set to 1, then the second (right-hand) instruction of the pair is modified by adding to it the (new) contents of the location specified by the address-part of the first instruction. This facility enables any location in main memory to be used as an index register. Address-modification is performed without loss of speed.

The main memory consists of 1024 words of 33 bits each (thus equivalent to about 4K bytes in modern terminology). The first four words contain fixed instructions, used for

initialising (boot-up). The memory technology is ferrite cores.

Data representation. Numbers are stored as two's complement fixed-point fractions, in the range $-1 \leq n < +1$. No floating-point hardware is provided on the 802.

The basic character size is 5 bits, corresponding to Elliott's version of the 5-bit teleprinter code. Since a particular 5-bit pattern could either be interpreted in *Figure shift* or in *Letter Shift*, more than 2^5 (but less than 2^6) symbols can be represented in Elliott 802 Telecode. They include:

the 26 upper-case letters A to Z;
the numerals 0 to 9;
16 punctuation marks, arithmetic operators and special symbols;
3 print-mechanism controls (carriage return, linefeed, space);
select *Figure shift*;
select *Letter shift*.
blank.

Physical characteristics. The 802 requires a single-phase, 2kVA, ac supply – thus putting it in the class of ordinary domestic appliances in terms of power consumption. The AC mains input is made to trickle-charge a large battery, thus providing (i) some degree of isolation from mains voltage transients, and (ii) time for the computer to perform an orderly shut-down in the event of mains failure. The CPU, main memory, I/O circuitry and control console are together housed in an L-shaped desk, the longer arm of which is 223 cms long and the shorter arm is 159 cms. The power supply is in a separate unit measuring 84 cms by 76 cms. All units are 81 cms high. The 802 can claim to have been one of the UK's first minicomputers.

Input/output equipment. When first introduced, paper tape was the primary medium for input and output to/from the 802. Other I/O devices were added later (see also section E3X3). When used for on-line control purposes, the 802 was often configured to have three logical Input channels and

two logical Output channels. In a typical process control application, the channels might be assigned as follows (see ref. 6):

Input 1: handkeys on operator's console.
Input 2: switchable between: (a) many A/D converters providing digitised data from the process being controlled; (b) a real-time digital clock; (c) other paper tape equipment.
Input 3: programmers' paper tape reader.
Output 1: programmers' paper tape punch.
Output 2: switchable between: (a) many control signals to the process being monitored; (b) a visual display of plant performance; (c) typewriter or other printing device.

803 systems architecture.

The 803A was an 802 but upgraded in the following respects: fully-transistorised core store drivers and logic core reset drivers (now that power transistors had arrived in the marketplace); word-length increased to 39 bits so that a larger main memory could be specified by more address-bits per instruction; floating-point hardware; L-shaped desk

format replaced by a cabinet-based physical appearance. The 803A was slower than the 802 because the 803A still retained the circulating *Operations register* which had to be increased to 120 bits to accommodate the increased word length.

In order to increase operating speed, the 803B was re-designed to incorporate several separate nickel delay lines in place of the main *Operations Register*. In the 803, SCR, IR and ACC were each held in separate delay lines. The machine's rhythm was changed to a two-beat *fetch/execute* sequence. Each beat was 48 digit-periods, so that simple instructions took (2 x 48 x 6) = 576 microseconds. (For timing of other 803B instructions, see section E3X3).

With a word length of 39 bits, integers are represented as before. Floating-point hardware is provided. Floating-point numbers are represented in standardised form as: $a.2^b$, the 803B holding the mantissa, a, and exponent, b, as:

$-1 \leq a < -\frac{1}{2}$, or $a = 0$, or $\frac{1}{2} \leq a < 1$,

and: $-256 \leq b < 256$

Actually, in the internal floating-point representation the exponent B is held as (b + 256), ie B is in the range 0 to 512. This avoids the necessity to sign the exponent and results in a more convenient representation of floating-point zero (for which a = 0 and B = 0).

Backing storage for the 803B is provided by reels of magnetic film. This is similar to the one-inch magnetic tape provided by many computer manufacturers except that, in the Elliott case, the backing for the magnetic coating was 35mm celluloid ciné film stock with sprocket holes. Each 1,000-foot reel of magnetic film could store over 250,000 39-bit words. Compared with conventional magnetic tape decks, access speeds for Elliott film handlers were slow.

The daughter board construction was not used in the Elliott 803B computer, the cores, transistors (OC84) and resistors being mounted directly onto the main board which measured about 10" x 17". There was no 'master/slave' relationship between cores in the 803B computer. The cores were called "alpha" and "beta", depending on the phase of the trigger pulse that reset them. Each board (which had cores on it) had a pair of OC23 power transistors to generate the two trigger pulses for the cores on that board. Alpha and Beta cores had equal status, with logic functions being performed by both phase cores. A minimum 803B installation consumed 3.5Kw of power.

503 systems architecture.

The random-access memory is in two parts: a fast section of 8K words (40K bytes) and a backing section of up to 128K words in units of 16K words. Both types of memory consist of ferrite core stores. The cycle-time of the fast section is 3.6 microseconds. Access to the backing section is normally via autonomous transfers to the fast section at 15,800 words/sec., though individual words can be accessed from backing store in 71.9 microseconds.

Bulk storage is provided by half-inch, 7-track magnetic tape decks, of which the 503 can have a maximum of eight. The instantaneous transfer rate is 42,000 chars/sec. Variable length records are employed. With a normal packing density of 556 characters/inch and tape-reels of up to 2400 feet in length and records of 128 words, one reel can typically

store about 1,730,000 words in total. For this record-length, the effective maximum transfer rate is about 35 records/second.

The fastest Elliott 503 instructions (group 4 jumps) take 4.5 microseconds. The fixed-point operations on the 503 take respectively 7.2 microseconds (group 0), 8.4 microseconds (groups 1 and 2) and 8.7 microseconds (group 3). Floating-point add/subtract takes 13.2 to 30.3 (average 18) microseconds. The figures for floating-point multiply and divide are respectively: 31.5 to 41.1 (average 33.8) microseconds and 60.9 to 61.8 microseconds. If the instruction preceding a B-line is in groups 4 to 7 then its time is increased by 3.6 microseconds. Details of the instruction set (order code) are given in section E3X4.