

Personal recollections of the Computer Laboratory in the Electrical Engineering Department of the University of Manchester, 1955 to 1964.

By David Aspinall.
December 2012.

In the autumn of 1955 Tom Kilburn recruited several new post graduate research students. I was one of them with an NRDC Scholarship of £300 pa. At that time the Meg computer had finished being the prototype for the Ferranti Mercury and was being used by Tom and Dai Edwards to generate waveforms for their experiments on magnetic core storage technology. Mike Lanigan joined them in this investigation. Eric Dunstan had joined the group earlier and was investigating Magnetic Drum recording technology. Dick Grimdsdale was making final touches to his computer based on Point Contact transistors, later moving on to investigate Pattern Recognition with Frank Sumner. Pete Hoffman and Roy Hayes were joined by Dennis Jeffries in an investigation into an Electroluminescent Panel as a computer graphics output device.

For my part, I was to investigate the design of logic circuits using the newly available Junction Transistor. It soon became apparent that such Diode-Transistor circuits were possible, if a little slow. Meanwhile Mike Mclean was working towards the end of his PhD project to investigate a circuit based on a magnetic core to provide the AND-OR-INVERT operations, followed by a transistor as amplifier. I began working with him and together we investigated the Residue Number system of number representation in a computer. Whilst it offered some advantages in the implementation of arithmetic functions, it was clearly too cumbersome for use in a general purpose computer. Nevertheless we built an arithmetic unit, based on the residue number system, using the Core-Transistor logic circuits. It proved the point. Mike got his PhD and I was awarded an Msc. (Paper published in March 1958).

At this point, in 1956, Pete Hoffman invited me to join his group to investigate Thin Magnetic Films as the next technology for the main memory. The work involved evaporating mixtures of iron and nickel onto narrow glass tubes in a vacuum, then devising test equipment to establish the resultant magnetic properties and their switching times between magnetic states. Pete and I invented a unique method of exploiting these thin films. It was going to be a long haul to see the idea turn into a practical memory. Indeed, thin film technology did not become available until the early 1970s.

Late in 1957 Tom invited me to rejoin his group and, under Dai's supervision, work on an idea he had for a high speed carry circuit. I left Pete and a new post-graduate student, John Turner, to continue with the thin film work. Dai and I quickly demonstrated that a number of transistors in series could all be turned "on" as a series of closed switches and propagate a pulse at almost the speed of light. It took several months of experiments to find the best transistor for the job. Also to design a circuit to enable rapid switching to make and break the transistor switches in the carry path. We were able to present the results at a specialist discussion meeting on New Digital-Computer Techniques at the IEE in London in February 1959. Later, Ianto Warburton of Ferranti and I developed a circuit to copy the carry signal from the carry path into the standard OC 160 logic circuit.

Some time in 1958 Tom asked me to join him in his room where I helped him to tape a large sheet of paper onto his desk. This was to become the "blue print" for MUSE. My role from there on was to sit by Tom's side, to act as his sounding board, while he drew the data paths of the machine on the blue print and the timing diagrams, for control, on squared paper. Over several months I witnessed the solution of many problems which were being presented to the designers of a computer for the first time. Most of the innovations were extant on the Blue Print by the end of 1958 and it was with some surprise that we heard them described by Christopher Strachey at the February 1959 IEE meeting in London.

For my part, I was concerned about the performance of the floating point arithmetic unit. We had a very fast Adder Subtractor using the high speed carry path. (In recent American literature this is referred to as "The Manchester Carry Chain"). We needed high speed Multiplication and, to a lesser extent, Division circuits. Following a lead presented in papers by Tocher & Lehman, I was able to overcome the difficulties of the Tocher/Lehman systems by examining the digits of the multiplier three at a time. By providing the adder with multiplicand D and also $2D$ together with an initial computation of $D+2D$ to form, and store, $3D$, it was possible to multiply 39 bit numbers in one plus thirteen iterations of the Adder/Subtractor, rather than thirty nine. By this means we had both an adder and a multiplication system (each patented) which enabled Atlas to compete with the best in the world at that time. (A side effect was the use of octal number representation in describing features of the machine which fitted nicely with the final choice of the character length of six bits). We were unable to come up with a satisfactory dedicated hardware solution to the Division operation and had to resort to Extracodes in the fixed store.

Meanwhile I was laying out printed circuit boards for the adder and assisting Yao Chen in building the MUSE prototype on Dexion frames. On 30th. January 1960 we were able to demonstrate the successful operation of the 24-bit adder in the B-unit prototype. A paper, describing the Adder, was published by the IEE in November 1960. It also hinted at the octal multiplier circuit to be used in Atlas.

Sometime in 1960 the final design phase of Atlas was begun. We had been collaborating with Ianto Warburton and Gordon Haley, of Ferranti, for some time and they, with Yao Chen, were to be responsible for the central control and arithmetic unit. Eric Dunstan became responsible for the Drum system and Mike Lanigan, with Frank Sumner, became involved with the main Core Store and the Page Address Registers of the "One Level Store". Dick Grimsdale was assisted by Keith Bowden to complete the Fixed Store. I was handed the responsibility for the Atlas Magnetic Tape System which, up to this point, had hardly been considered by the main group. David Howarth, of Ferranti London, became responsible for the Supervisor software.

The choice of tape deck was to be made by Ferranti and was to serve both the Atlas and Orion computers. Reg Claber, of Ferranti, became responsible for looking after the interests of the Atlas computer. He proved to be an invaluable aide in the specification of the system performance. Ferranti sub-contracted the detailed design of the tape deck to Ferranti Edinburgh. This involved many meetings in Manchester and Edinburgh, once the choice of the Ampex TM2 Tape Deck had been made. The Ferranti-owned aeroplane came in useful for several trips to Bill Alexander of the Ferranti Craigroyston Laboratories, on the shores of the Firth of Forth.

The specification for the Atlas magnetic tape system became:- Up to 16 tape decks, each with a transfer rate of 90,000 characters per second, connected as eight dedicated channels each channel to be switched between two tape decks. Simultaneous transfers between all eight channels and the main store were to be allowed. The system had to cope with blocks of a fixed length of 512 words on pre-addressed tape. The pre-addressing was to be carried out by hardware in the Tape Co-ordinator, controlled by software in the Supervisor System. This part of the Supervisor together with the test program for commissioning and maintenance were written by Ron Lane from Ferranti. The complexity of the Tape Co-ordinator grew to require three computer bays. These included a core store of 16 words of 50 bits (8 Characters of six bits plus 2 bits of parity), cycle time of one microsecond, plus a 24 bit adder to compute the check sum, per block on tape. A novel read amplifier, to a design suggested by Tom, was included in the co-ordinator.

Orion planned to place the circuit for switching between decks and channels under the floor boards. We left this problem to them, as we were only going to have eight decks on the Atlas at the University of Manchester.

Detailed design of the Co-ordinator was passed over to the factory at West Gorton by July 1961. Leaving Reg Claber, Ron Lane and a commissioning engineer to hold the fort, I went and stayed at home, for six weeks, to finish my PhD Thesis. In September Dai agreed that it was sufficient and my external examiner, Bill Renwick, from Cambridge, was content. I was admitted to the degree in December 1961. By this time we had begun to commission the magnetic tape system.

The priority at the time was the commissioning of the Atlas Central Processing Unit, The Page Address System, The Core Stores, The Fixed Store and the Drum to enable development of the Supervisor and compilers. The result was that the Tape System was allowed only limited access to the rest of the machine. This often meant that we were given the total system on Friday night until Monday morning or earlier if the rest of the machine failed. If it failed before Saturday night we would ring Yao Chen at home for him to come in and fix it. Thankfully he was willing to do this on more than one occasion.

By the late spring of 1962 we had commissioned one channel, with its tape deck, and had addressed several tapes. We were open for business and the programmers, including Tony Brooker's team, began to use Magnetic rather than Paper Tape. We rose up the list of priorities and began to gain access to the machine to finish the commissioning of the remaining channels.

In the run-up to the 7th December it became essential to give David Howarth as much time on the machine as possible. This included the Tape System which was still subject to failures as the dry joints and faulty OC160s were weeded out. We tried to get it working again for David as quickly as possible. David never seemed to sleep. On one occasion he arrived from London on Monday morning. I resolved to stay and be available as long as he was awake. By Wednesday morning I was exhausted and went home to sleep, leaving him still on the job.

On the day of the official Atlas switch on, David was still debugging the supervisor. After the actual switch on, by Sir John Cockcroft, the main party moved away and we were able to actuate all eight Tape Decks and perform simultaneous transfers between them and the

main core store for the first time: Job done. ... except it wasn't! There followed several months of intermittent faults, due mainly to dry joints and faulty transistors. Whereas the commissioning engineer was fully proficient, the maintenance engineers were still learning the ropes. I backed them up 24 hours a day from Monday to Friday. As the MTBF increased from hours, to days, to weeks, my involvement diminished. One morning, on reading the log, I found the maintenance engineers had cleared a difficult fault during the night, without getting me out of bed. From then on they were on their own. The Tape System had become an essential part of Atlas. It enabled the Ferranti Computer Service to offer Data Processing facilities.

Also data from the 1961 Census for the North West of England was copied from punched cards onto Magnetic Tape, for the benefit of the Department of Economics at Manchester University.

A testimonial on the magnetic tape system from Bob Hopgood (of Harwell) was well received recently, when he made this comparison of the Atlas and IBM STRETCH computing facilities:

"My codes were pretty similar in performance but on large calculations where intermediate results had to be stored on magnetic tape, Atlas was significantly faster due to the Ampex tape decks. I think on an untuned general purpose workload Atlas was faster and if the code was tuned to STRETCH it would be faster. In conclusion, I would say that in 1962 Atlas was reckoned to be the world's most powerful general-purpose computer".

There were two other pieces of Atlas hardware that I became involved with. The first was an attempt to provide high speed text out put via a CRT screen photographed onto 35mm film, developed and magnified to produce hard copy. The equipment was designed and constructed by Data Laboratories, a new company, in London. I was asked to interface it to Atlas. Fortunately the Analex Line printer arrived before we had spent much time on the project.

Next there was a Digital Modem developed by British Telecommunications Research Ltd., a branch of Plessey. We developed the interface to Atlas plus the necessary software and performed tests on a data link, over telephone lines, to UMIST.

In 1964 I transferred my lectureship from The Department of Electrical Engineering to be a founder member of The Department of Computer Science at The University of Manchester.

It is interesting to consider what became of the members of the Atlas team who were in the group in 1955. The most senior member, at that time, was Tom Kilburn, Reader in Electrical Engineering.

Except for Eric Dunstan, who went to California to earn his fortune, we all became professors:

Tom Kilburn	CS Manchester
Dai Edwards	CS Manchester
Dick Grimsdale	EE Sussex
Tony Brooker	CS Essex

Mike Lanigan EE Kent at Canterbury
David Aspinall EE Wales at Swansea, then Comp at UMIST

Other members of the Atlas team, prior to August 1961, who became professors, were as follows.

Frank Sumner CS Manchester
Derrick Morris CS Manchester, then Comp. at UMIST
Jeff Rohl Comp. at UMIST, then CS Perth Australia
Bruce Payne CS Tasmania Australia
Keith Bowden CS Essex