The Atlas computer in the 1960s and early 1970s: a personal reminiscence

John S. Reid

Introduction

I was a long-term user of the Chilton Atlas computer in the second half of the 1960s and early 1970s. The Chilton Atlas Laboratory adjacent to the Rutherford High Energy Physics Laboratory was a national facility, particularly used by university researchers who needed a larger and faster machine than local computing could provide. When it was inaugurated in 1966 if was one of the largest and most innovative computers in the world – a supercomputer before the word had been coined. This note includes a few reminiscences and some technical details of the machine, extracted from the archival pages of http://www.chilton-computing.org.uk/. It accompanies items in the Natural Philosophy Collection that include a collection of pictures showing the operations room and other aspects of the Atlas computer.

Among the relevant artefacts in our collection are a manual card punch (ABDNP:200887a), program and data cards (ABDNP:201997a), a programming pad, assorted unpunched IBM cards (ABDNP:200887b), a card-sorter (ABDNP:201485a) not connected with Atlas, a postal box for sending cards to the Atlas computer facility (ABDNP:200843a) and Atlas lineprinter output (ABDNP:201998a), 1-inch Ampex tapes (ABDNP:202002) and ½-inch IBM tapes (ABDNP:202003a), ferrite core memory (ABDNP:201920a), diode memory from a card punch (ABDNP:201352a), a mainframe disk storage stack (ABDN:201940a), paper tape punch (ABDNP:201027a), paper tape reader (ABDNP:200848a), 5-hole (ABDNP:201996a) and 8-hole paper tapes (ABDNP:200845a) and manual paper tape editor (ABDNP:200859b), and various circuit boards illustrating the electronic technology of the 60s. These items together give a flavour of some of the academic computing in the 1960s, the first decade in which computing moved from a few pioneering centres to the University system at large, albeit still mainly taken up by only a small number of disciplines such as Physics, Physical Chemistry, Mathematics, Engineering and the rising area of Computer Science. The biological and social sciences were comparatively slow adopters.

Over 80 images of the Atlas laboratory and its computing systems referred to here are contained on a CD (ABDNP:201993a) and a DVD (ABDNP:201993b) to accompany this account. They have been sourced with permission of the Rutherford Appleton Laboratory (RAL) and the Science and Technology Facilities Council (STFC) from images on http://www.chilton-computing.org.uk/ and from RAL's image library. These images may be viewed on our collection web-site in due course but anyone wishing to re-use these images should obtain them from RAL (STFC) and accompany them with the required acknowledgement. Fuller image details including caption information are listed in the catalogue entries for ABDNP:201993a and ABDNP:201993b.

Background

I was part of the Diffuse Scattering Group of Dr Tom Smith. Our local university computer was an Elliott 803B, purchased in the early 1960s. I remember it as having 16 K of storage (wrongly, for I see from contemporary documentation that it had 8K of storage) and could be programmed in *machine code* (individual hexadecimal instructions), *autocode*, a simple symbolic representation of machine code, and *Algol 60* (a high level language). Our collection includes a cabinet of 5-hole tapes of programs, dumps and data for the 803 (ABDNP:201996a). I wrote my first Algol program in 1964 to calculate the 1-phonon x-ray scattering predicted for the alkali halides, NaCl in particular. This was a comparatively simple calculation, achievable with care for a single scattering direction using a mechanical hand calculator but something that would have been very tedious to repeat for many directions, which the computer could do quickly. Our computing requirements soon expanded to working out 2-phonon scattering, a task beyond a hand calculator and, later, other scattering processes. These required input from a lattice dynamical model that predicted phonon characteristics, opening up a second area of computation. In addition, I spent about 6 months altogether at Harwell from the summer of 1966 carrying out neutron scattering experiments and analyses on NaBr aimed at measuring experimentally phonon energies for this material, which was a crystal of current theoretical interest, was simple enough to model and was one of the alkali halide family that had been least explored. As far as I know the results still represent the definitive values for NaBr, some 40 years on [J. S. Reid, T. Smith and W. J. L. Buyers *Phonon Frequencies in NaBr* Physical Review B, vol. 1, pp 1833 - 1844 (1970)]. This activity got me into using FORTRAN for analysing the NaBr scattering data, the Harwell and Atlas computing facilities and moving to FORTRAN for all my computing. The Harwell site was adjacent to the Atlas Lab, some 15 miles south of Oxford.

Our group had therefore 3 strands to its computing activities: calculating predicted phonon scattering; calculating phonon characteristics (namely particular properties of a crystalline solid) from a range of lattice dynamical models (namely, mathematical descriptions of how the atoms in a solid vibrate) and finally converting experimental measurements of X-ray and neutron scattering into useful information about the scattering material. Now I come to think of it, there was a 4^{th} strand: calculating other properties of solids that depend on the phonon characteristics once these had been evaluated from the models. Others in the group involved in computing at this time were John Pirie (particularly lattice dynamical models and anharmonic scattering using Atlas), Stuart Melvin (various aspects of lattice dynamics using the Atlas machine), Gordon Peterson (scattering from aluminium using the Elliott 803 and its successor) and in the early years using the Elliott 803 Bill Buyers, who moved to Canada later in the mid 1960s but continued to collaborate on the NaBr project using the Chalk River computer.

A significant amount of this effort required computing facilities beyond the capacity of the Elliott 803, though we did stretch this to its limit. Our group were frequent all-night users. There was no operator. One turned up at 10 pm, edited the 5-hole tapes of program or data if necessary (though that could be done beforehand) and operated the machine. Translating an Algol progam took some time and if all went well one left about midnight with the machine burbling to itself and the calculation proceeding until it stopped or was stopped between 8 and 9 in the morning. The output was paper-tape only, which had to be printed in an adjacent room. Gordon Peterson remembered recently "I still have painful memories of the static electricity that used to discharge into us when we picked up the metal bucket that we collected the paper tape in. A real example of Faraday's ice pail." Paper tapes were almost un-editable. One could add a few holes at the front with a hand-punch if some characters were missing, for example, but mainly any edits to a program required that the complete tape was re-made. Looking though some old print-out, it looks as if the Atlas could read Elliott Algol tapes but most Altas computing was done in FORTRAN using punched cards. What about the Atlas Computer?

The Atlas computer in brief

The Chilton Atlas was one of 6 Atlas computers built by Ferranti and developed in conjunction with the University of Manchester and, later, the University of Cambridge. The Atlas 1 developed in the early 1960s could claim to be the fastest and most innovative computer in the world, or certainly on a par with the 'IBM 7030: stretch' that might dispute the claim. The Chilton Atlas's remit was to provide a national computing facility for research in science and engineering. It came into full operation in 1966 and was switched off in March 1973. Facts in this section can be found along with fuller details in http://www.chilton-computing.org.uk/acl/home.htm.

Atlas was a full multi-tasking machine, allowing many user programs to be in the system at once and, indeed, many to be executing at once. This was unlike the University's Elliott 803, or indeed the first generation of micro-computers.

The Atlas word length was 48 bits. Atlas internal storage was quite complicated, being organised hierarchically among several different pieces of hardware. In short it had 48 K (words) of ferrite core, access time 2 μs; 8 K fixed (read-only) store, access time 0.8 μs; 16 K working store, access time 2 μs and 96 K magnetic drum store, access time up to 12 ms depending on where the information was on the drum. One of the accompanying photographs shows Atlas ferrite core. In 1968 a 16 million word magnetic disk was added. The Atlas was built using discrete component transistor technology.

Atlas computer hardware, below stairs

Storage and engineers' console

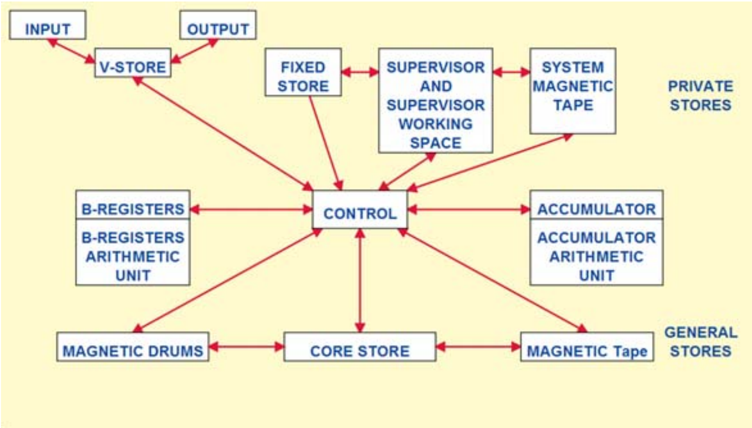

Architecture of Atlas computer

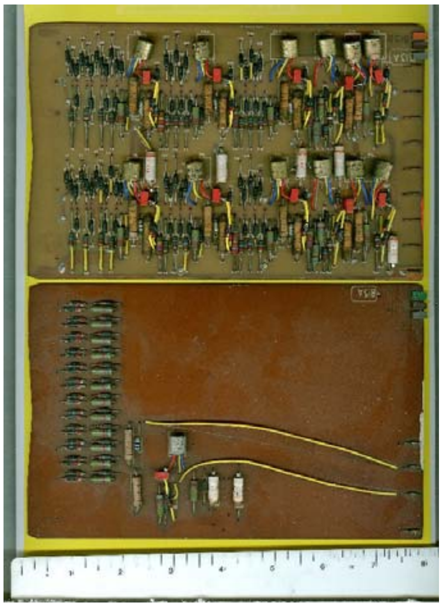

Example circuit board

Atlas B registers

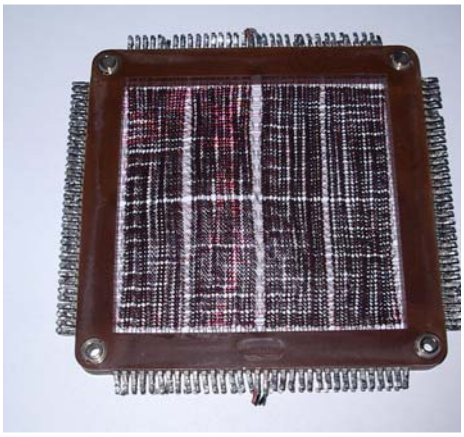

Atlas ferrite core module

When each word represented a floating point number, 8 bits (including a sign bit) were used for the exponent and 40 (including a sign bit) for the manitissa. In round figures, floating point addition in the arithmetic unit took just under 2 µs, multiplication about 6 µs and division a variable time around 20 µs. Scientific calculations were therefore by default done to higher precision than in the 32 bit word machines that became standard over the following 3 decades. For comparison times, the Elliott 803 floating point addition for its 39 bit words took about 860 µs and multiplication almost 5000 µs, i.e. 5 ms. So for this kind of arithmetic the Atlas was over 500 times faster. An 8-hour arithmetically intensive run on the Elliott 803 could be done in less than a minute on Atlas. That's the theory. In practice many other issues influenced the time taken and I'm not sure if we ever ran any comparisons with the same program on both machines. Also in practice our Elliott didn't have the storage to hold the information needed by the Atlas programs.

Atlas operations room soon after installation

Atlas operations room in action

The main tape system at the back of the Atlas operations room stands out in the photographs. It used 16 1-inch Ampex magnetic tape drives, each taller than the average person. One tape (of 3600 feet, 1100 m) was divided into about five thousand 512 word blocks, a word having 48 bits. There were a standard 16 tracks across the width of a tape, 12 for the data, 2 for timing and 2 as checksum bits. Hence each word occupied 4 data stripes of the tape. Putting this together gave each tape a capacity of a little over 20 million words, or about 1 Gbit of information. Tapes could be read forwards or backwards at 120 inches per second, giving an effective data transfer rate of 64 K words per second, allowing for block markers on the tape and gaps between blocks. That makes it 6 minutes to read a whole tape, end to end, if that were necessary. There were also two, smaller, ½-inch IBM tape drives.

Most user input was provided by 80-column punched cards, read at a maximum speed of 600 cards per minute; or 5-, 7- or 8-hole paper-tape, read at 300 or 1000 characters per second. Output was provided on paper-tape by two Teletype punches, maximum speed 110 characters per second, two card punches, 100 cards per minute, or two lineprinters with 120 characters per line printing at 1000 lines per minute.

In pictures of the computer room a silence reigns supreme but in reality the chatter of lineprinters, the zip of tape readers, the zizz of the tape punches, the whirr of tape drives, people moving cards, tapes and paper or talking over the noise of air conditioning created a continuous sound to accompany the continual activity. The glass partition between operations room and user hall shielded the users from this noise and also shielded the operators from the chatter of users and the data-preparation punches in the user hall.

Apart from a large flat-bed plotter, one further device made the Atlas computer very special: the SC4020 graphical output device by Stromberg-Carlson. This used a very fast cathode ray tube that could generate characters or graphics. The output was both photographed onto film and made available as hard copy (by a light-beam splitter that directed the image of the tube face onto both a film camera and photographic paper). Both the film and hardcopy had to be processed by the usual photographic chemistry. Exceptional speed was achieved by an arrangement in which the electron beam was first directed through an internal mask in the tube cut with all the letters of the alphabet and many more shapes. The shaped electron beam was then directed to the correct place on the tube face to make up part of the image. In this way a whole character was imaged at once instead of having to be made up by scanning the electron beam across the screen to draw out the shape of the character. Users didn't see the

Postal reception at Atlas Laboratory

Card boxes on trolley at Atlas postal reception

machine, requested output just appeared. The photographs show that the SC4020 occupied 5 tall cabinets adjacent to a tape drive. From 1969 it was run from a dedicated off-line PDP15 computer. In addition the photographic output had to be chemically developed.

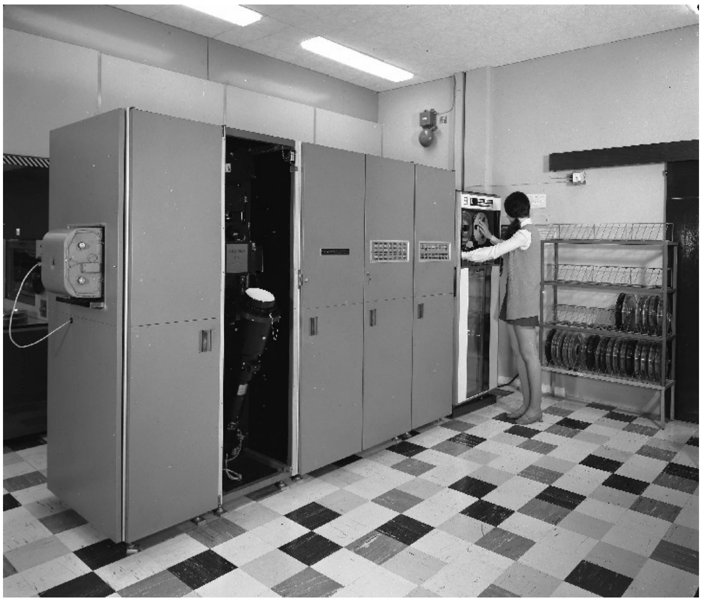

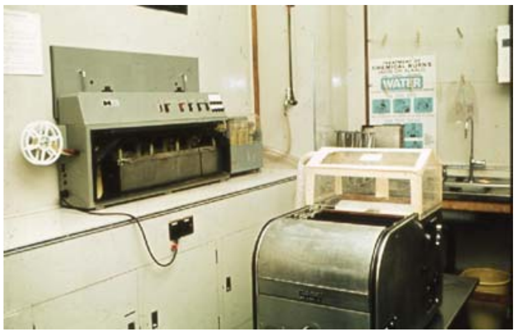

Tape loading the SC4020, with image tube visible

Dark-room processing for plotter film

An audit of a typical week's work in the Atlas lab in 1967 that gives a flavour of the throughput showed that the Atlas ran 2,500 complete jobs; read in a million cards and 30 miles of paper tape; printed 2 million lines of output; punched 50,000 cards and handled 1,200 reels of magnetic tape. Computers in the 60s were physically big and so were their peripherals, even though storage by modern standards was tiny and speeds were slow. They did, though, get through a lot of work.

A few memories

As far as I remember, my use of the Atlas covered the years 1967 – 1972. (In 1973 the facility was closed down since its replacement, an ICL 1906A, installed in a neighbouring new block, was fully in commission). By the late 60s the University had updated the Elliott 803 to a substantially larger machine (an ICL System 4/50) that ran FORTRAN programs from card and, after some user pressure, wrote user output to half-inch magnetic tape drives. Diffuse X-ray scattering calculations in the 1960s used to evaluate 1-phonon scattering, 2-phonon scattering and, with some sweeping approximations, 3-phonon scattering. One result that I developed that has stood the test of time was to show how one can evaluate all such phonon processes at once, not just ending in 3-phonon processes but taking the series to its limit – the so-called complete multiphonon scattering. The solution is at first sight too computationally intensive, involving sums over 6 spatial dimensions, but computing in the 60s had generated the ethos that there were often clever ways of tackling problems when resources seemed very limited, as they were computationally both in terms of storage and speed. The program I developed for the Atlas computer showed that multiphonon calculations were possible as a matter of routine when such a machine is available [J. S. Reid and T. Smith *The Calculation of Multiphonon X-ray Scattering* Journal of Physics C, vol. 3, pp 1513 - 1526 (1970)]. As far as I know, no-one has improved on the method in the succeeding 40 years.

The concept of 'structured programming' hadn't been widely circulated in the late 60s but a typical FORTRAN program consisted of a modest length of main routine that called subroutines to execute specific elements of the calculation. These in turn could call further subroutines, and so on. Writing a program involved first covering a lot of sheets of white paper with blue ink, deciding how the whole calculation could be broken into subroutines and then sketching the content of individual routines. When the details of a routine seemed clear enough, the necessary FORTRAN statements were printed by hand onto green FORTRAN coding forms, an initial supply of which came from the Harwell lab and then we had additional pads of forms printed. These forms were ephemeral items, being thrown in the bin when the cards had been punched.

I think that both the Atlas laboratory, and the University in the 1970s, had a data preparation section that one could send the forms to and receive the equivalent set of 80-column cards in return, each card with the text printed along the top. (At the time, few academics typed their own letters, reports or papers, for in each Department of moderate size there was a Secretary and several typists. Typing the punched cards, though, needed a special card-punching desk, usually located only in computer centres or certain admin sections). The Atlas photographs show cards being prepared from coding forms using a card-punch. Some card punches printed the text that was coded on the card; others didn't, usually the punches that the computer used for output. Card output generated by computer was normally destined to be data for another program and hence often didn't have to be easily readable by eye. However, if one dropped the cards, or even let one or two slip, getting them back in the right order was a challenge if they hadn't their text or numbers printed on them. One or two card punches were 'interpreting', meaning that one could feed punched cards in and the machine would slide each card along, interpreting the holes and typing the corresponding text along the top as the card passed through and out the other side. Another aid to sorting program cards was that FORTRAN statements used only columns 1-72. Columns 73-80 could be used for a running number, effectively a program line number, though data-preparation staff seldom typed anything into these columns.

Card punch room in the early days

Card punch room in 1969

Experience soon showed that if the set of routines that constituted the program was immediately assembled and run then the program would crash, due to such causes as a failure in the author's logic, coding errors, typing errors, data error, a special case that one hadn't treated appropriately, and so on. There seemed to be many ways for a program to fail. I recall that it took me weeks to track down one error in the multiphonon program where the digit 1 had been typed as the letter l in a context where both did not upset the syntax of the program. (In some fonts such as Time Roman, the letter and digit look identical but the computer doesn't read fonts but interprets characters and the characters are different. The digit 1 typed as the letter l was hard to spot on a printout). So routines were usually tested individually, writing special test programs for each routine. This doesn't eliminate all mistakes but helps a lot.

The hopefully working program was sent to Atlas by post in a special box made for computer cards. Photographs show the postal reception and dispatch at Atlas. Our programs needed data and the data in our case was generated by another program. The output of this program was on cards at first and of course the program generating the cards could fail or, worse still, run but produce erroneous numbers on its output cards. Later we used magnetic tape for the intermediate numbers. There was a time in the development stage when the best plan was to go to Atlas to speed the turnaround. I had an old Austin 8 in the mid 1960s, a car of pre-war design that was slow, burnt a lot of oil but was very reliable. With a foot of snow on the bonnet it would start almost immediately, if not with the starter then with the starting handle. 500 miles in a slow car that was poor at overtaking even slower lorries and buses, on routes that lacked many of the dual carriageways we now have, wasn't a great experience but needs must when the science drives.

Atlas library

The arrangement at the Atlas lab for visiting researchers was that they would be put up in a local pub, usually near Abingdon. The landlord would be warned that we'd sleep all day. I rose in the late afternoon and headed for the lab in the early evening with boxes of cards, printout, magnetic tape if necessary and accompanying notes. The first stop, I think, was the evening canteen. The chef had a large, lightly oiled, hotplate. One selected a steak, beef or gammon, and it was braised to spec as one stood listening to the hissing crackle, watching its juices ooze and smelling the aroma with anticipation. The hassle of driving 500 miles suddenly seemed worth it. What a simple and effective ploy. Who wouldn't want to come in to work and spend about 12 hours right through the night in a good mood without complaint after such an evening meal? It never ceased to be a pleasure to go to the Atlas lab.

It helped that the staff were sympathetic and well organised, as was the whole lab. The operations room and working area were upstairs. The metal cabinets that housed the heart of the machine were downstairs, supervised by a team of engineers, unseen by most users. Visitors working the night-shift got job priority on the Atlas. Hand in the deck of cards to the operations reception at the glass partition and it went straight to the front of the queue. One could watch the card-reader swallowing cards, the Ampex tapes jerkily spinning, if not for my job for some else's, the lineprinters chattering out pages. If the job failed then the error was hopefully soon spotted on the diagnostic print-out that accompanied the program output; new code cards or a change of data could be punched on one of the many card punches in the hall in front of the operations room and the deck re-assembled and the job re-tried. If the job was successful then the next development of the program could be attempted.

I can't remember exactly what the gain from being present at the Atlas lab was but a 3-night visit was likely as productive as several months of postage operation from Aberdeen. I finished at about eight in the morning and returned to the pub for breakfast. The journey home at the end was tedious in the Austin 8 but the reliability of the Atlas machine was impressive (97% uptime was its lifetime figure) so I was never going home empty-handed. In the late 1960s the Austin reached the end of its life and to the amazement of my flatmates I exchanged it for a Jaguar XK140, famed in the 1950s for its exceptional performance in the Le Mans 24-hour race. No more driving for miles behind slow lorries or peering through a mucky windscreen at the back of a bus for 10 minutes on end. After some time, car journeys gave way to the plane to Heathrow and I was met by a laboratory car. This seemed luxury treatment to me as a young researcher but it was efficient for me and a minor cost for the lab, given that they already had in-house cars with drivers for their own use.

The usefulness of the Atlas computer steered the thrust of my research at the time away from laboratory experiments into computational physics. I developed particular aspects of the theory of anharmonic scattering and with Dr John Pirie we used Atlas to calculate this effect [e.g. J. D. Pirie, J. S. Reid and T. Smith *Lowest-order Anharmonic Thermal Scattering of X-rays* Journal of Physics C, vol. 4, pp 289 - 297 (1971)]; J. S. Reid *Higher Anharmonic Contributions to Diffuse X-ray Scattering* Physica Status Solidi (b), vol. 52, pp 291 - 301 (1972)]. Later computationally intensive work investigated the effect on X-ray scattering of dynamically deforming atoms [e.g. John S. Reid *Deformation Coupling of X-ray Scattering Factors* Journal of Physics C, vol. 7, pp 3444 - 3451 (1974)] and such work continued during the 1970s after the Atlas had been replaced. A small collection of original off-prints (ABDNP:201998b) of the papers mentioned here is included among the artefacts. As I recall, by the mid 1970s local computing facilities had improved very significantly, an ICL System 4/70 replaced the 4/50 and in 1977 that itself was replaced by a Honeywell 66/80; work that was initiated on Atlas could be run locally.

Computing facilities of the 1960s look so different to today's operations that those who haven't experienced them must wonder when looking at pictures of them what they were all about and what they could do. This piece is intended to give some facts and figures and a little of my experience of the times. One final figure about the Atlas computer: the fraction of CPU (central processing unit) usage taken up by users' program operation was 82%, by the operating system (the supervisor and scheduler) 12% and time spent idling for one reason or another, 6%. Do any modern computers come near to this user efficiency?

The Atlas Computing Laboratory was merged with the Rutherford Laboratory in 1975.

John S. Reid

V1.0 September 2010

V1.1 illustrated, July 2011

ICL 1906A, the Atlas Computer replacement in 1972