Celebrating 50 years of computing at the Rutherford Appleton Laboratory

In 1964 the Atlas Computer Laboratory began to provide a scientific computing service on one of the most powerful computers in the world, a Ferranti Atlas 1. For nearly a decade the machine supported scientists from all over the UK, and the Laboratory was at the forefront of developments in software and other enabling technologies aimed at making life more productive for its users. In the 1970s the Laboratory was absorbed into the neighbouring Rutherford Appleton Laboratory (RAL) on the Chilton campus, and in subsequent decades the computing facilities expanded enormously in scale and a wide range of other IT support and development activities were undertaken.

This celebratory event took place on the afternoon of Thursday 13th November 2014. Its purpose was to review the history of computing on the RAL campus over the past 50 years and then to hear about the computational challenges and opportunities facing today's scientific community. The event was attended by about 180 people ranging from members of the original Atlas Laboratory to today's staff and users. The talks were followed by a reception in an exhibition area where equipment and artefacts from previous decades were on display.

The speakers' slides are available at:
http://www.chilton-computing.org.uk/acl/technology/chilton50th/overview.htm .

A comprehensive history of computing on the RAL site from 1961 to 2000 can be found at:
http://www.chilton-computing.org.uk/ .

After brief introductory and welcoming talks by Juan Biccaregui, Head of the Data Division of the STFC Scientific Computing Department, and Andrew Taylor, Director of STFC National Laboratories, the following talks were presented. The emphasis in this article is on the history rather than the future, and it is difficult to do justice to the 'future' presentations in brief summaries. Readers are encouraged to look at the detailed 'future' presentations at the url given above.

The history

Computing at Chilton 1959-2000

Bob Hopgood, who played a major role in computing at Atlas and RAL from the early 1960s to the turn of the century, gave the first talk in which he faced the formidable challenge of covering this 41 year period in 30 minutes. In fact he started earlier than 1959 with a summary of the development of scientific computing in Britain from the mid-1950s when about 75% of all computing capacity available to scientists was owned by the UK Atomic Energy Authority. In 1961 AERE Harwell made the case for an extremely powerful machine and after much scientific and political debate it was decided that the machine would be a Ferranti Atlas, that it would be housed in a new laboratory outside the Harwell security fence and that its capacity would be shared between Harwell, universities and Government organisations and the National Institute for Research in Nuclear Science (NIRNS) which would be responsible for the facility's operation. Thus the Atlas Computer Laboratory came into being, its first Director being Jack Howlett.

Bob described the machine's architecture and parameters, noting that the UK could then claim to have the most powerful general purpose computer in the world and that the Atlas doubled the computational power available to the UK academic sector. Jack Howlett's enlightened use of his discretionary powers as Director fostered new areas of scientific computation. A major, vital, early in-house development was of a Fortran compiler and, in addition, the Brooker-Morris compiler–compiler was used to generate compilers to enable software written in other languages to be run on Atlas.

Providing good graphical output, whether on paper, film or microfiche, was an early priority. This was done via a Stromberg Carlson SC4020. A microdensitometer service based on an Optronics P1000 scanner was introduced in 1974 serving 30 university groups working mainly on protein crystallography. Interactive editing and job submission services were provided in the early 1970s through a Sigma 2 front-end computer which had independent access to a 16 Mword disc shared with the Atlas machine itself.

The SC4020 was replaced in 1975 by an FR80 microfilm recorder which was the most versatile and most accurate of the recorders then available. As well as being able to produce output to paper and film it could produce high density microfiche, so that 400 pages of output could be contained on a single fiche. Its ability to produce colour output was unrivalled at the time and the FR80 was to continue in service until 1985 when its single-frame output functions were replaced by specialist printers, plotters and a microfiche recorder. The multiframe functions were replaced in 1990 by a facility that produced video rather than film, and a virtual reality facility was added in 1997.

The main scientific disciplines using the Atlas were chemistry, physics, mathematics and theory and engineering. Projects included weather forecasting, satellite data processing, the development of a major quantum chemistry software suite ATMOL, X-Ray crystallography, modelling the strength of scaffold towers, M6 bridge construction, effluent in Southampton Water, and simulating hovertrains.

The Atlas machine was closed down at the end of March 1973, having achieved 97% up-time and having processed 836,000 jobs. Some of the equipment survives - five bays of the machine are housed at the National Museum of Scotland (and the Atlas console was on show at the exhibition accompanying this event). Considerable study and long negotiation went into possible successor machines. The eventual decision was to install an ICL 1906A and also to have access to 20% of the capacity of an IBM 360/195 that had recently been installed at RAL primarily for particle physics use. The 'Atlas' usage of the new facilities would be confined to SRC grant holders. Prior to the 360/195, RAL had a Ferranti Orion followed by an IBM 360/75.

In 1975 financial constraints within the SRC led to a decision to close the Atlas Laboratory as a stand-alone institute, and its staff resources were dispersed. Staff resources (but not in most cases the actual people) to support SRC Science Board activities transferred to Daresbury Laboratory, where in the late 1970s a Cray-1 vector processing service would be introduced. RAL took on the management of the ICL 1906A and related services and 55 staff-years were transferred to SRC's Engineering Board to set up a new Interactive Computing Facility (ICF) for university engineers. The ICL 1906A was eventually closed down in 1978. The IBM facilities at RAL were moved into the Atlas building which was renamed the Atlas Centre.

The following years saw progressive expansion of the large scale machines at Atlas. The IBM equipment went through many upgrades and replacements, driven mainly by the needs of experimental particle physics for data processing and monte carlo simulation, until by 1988 it had become an IBM 3090-200E with a vector facility.

In the meantime the Engineering Board's ICF was developing along the lines set out in a 1974 report on Engineering computing Requirements (the Rosenbrock report). The emphasis was on interactive computing, initially though enhancements to existing DEC 10 systems at Edinburgh and UMIST, installing Prime and GEC multiuser interactive systems at other universities and moving towards the provision of single users systems connected to local area networks through which they could access specialised facilities. Six special interest groups were set up to foster software developments and standards: artificial intelligence, digital and analogue circuit design, electromagnetics, finite elements, control engineering and computer aided architectural design.

The first commercially available single user systems were PERQs, designed by Three Rivers Corporation and manufactured and sold to SERC by ICL, and running a version of Unix designed jointly by ICL and RAL. These were the first personal computers in the UK. Later over 200 Sun single user systems were distributed to university users, and superfast workstations supplied by Stardent and AMT were also assessed.

When Atlas was first installed the transmission of input and output between the machine and its users was by physically transporting it by land or air. In the early 1970s radial networks began to appear from computer centres to their users and over the next decade a process of integration and standardisation took place which led eventually to the creation of the Joint Academic Network. One step along the way was the creation of SERCNet which brought together the separate RAL and Daresbury external communications systems, and had 186 computers attached by 1983. In 1977 there was a successful demonstrator project that linked for the first time SERCNet via UCL to ARPANET and other networks in the USA.

After the ICF programme came to an end the Laboratory managed and supported several SERC and DTI computing initiatives: Distributed Computing Systems (1977 – 1984), Engineering Support (1986 – 1995) the Alvey Programme (1983-1987) and the Transputer Initiative (1987-1992), all of which required RAL to support large numbers of projects across the UK academic engineering and computer science communities. Also, from the mid 1980s to the late 1990s, RAL played a leading role in a large number of ESPRIT projects in which SERC funding was enhanced by Esprit research funding.

Returning to RAL's large scale scientific computing activities, in 1985 the ABRC set up all-Research Councils Working Group to define future needs for high performance scientific computing. It recommended that a Cray X-MP/48 be installed at Chilton to be available to grant holders from all Research Councils. This machine was in service from 1987 to 1992, was replaced by a Cray Y-MP8128 which ran until 1996 to be replaced in turn by a Cray J-932 which ran until 1999. Also during this period IBM offered RAL a Joint Study Agreement whereby the company would upgrade the existing IBM 3090-200/1VF to a maximally configured 3090-600E/6VF (six processors each with a vector facility) and RAL would bring to the machine a set of numerically intensive prestigious projects from academics around the UK representing a wide variety of science and engineering applications and host a seminar at the end of the study on achievements that had been made.

Disciplines benefitting from these machines included global atmospheric and oceanographic modelling, lattice gauge theory, fluid dynamics, material science, plasma physics, quantum chemistry, mechanical design and structural integrity, astronautics, astronomy, drug design and protein crystallography.

By the early 1990s it was becoming possible to provide high performance computing facilities on powerful and relatively inexpensive modular systems. Processor farms for particle physics computing were soon providing much more computing capacity than had been available via mainframe technologies. Dedicated systems, initially using DEC equipment, provided for projects from other disciplines that required high performance but not necessarily at the level provided by national supercomputers. These included, in 1997, a Fujitsu VPP300, which had three of the most powerful vector processors then available, for use by NERC projects that needed high computational power concentrated in a small number of processors..

It was apparent from the early days of Atlas that storing, managing and providing access to large amounts of data would be an ongoing challenge. Initially the only available solutions were labour intensive: thousands of magnetic tapes being handled manually. Over the years more automated storage devices were installed such that, for example, by the mid1990s an IBM 3494 linear tape robot system could house up to 20 Tbytes of data. A major task in this period was to develop software systems that could provide access to the data store from different types of networked computers that was a future-proof as possible as technologies developed.

In the 1990s there was considerable activity at RAL relating to ERCIM and to the development of the World Wide Web. ERCIM is the European Research Consortium for Informatics and Mathematics and was founded by laboratories in Germany, France and the Netherlands. RAL joined the organisation in 1990 and by 1994 there were a further seven members. In 1990 RAL had installed one of the world's first web servers and was to participate in web standardisation activities through the activities of the World Wide Web Consortium (W3C). RAL provided graphics and XML support for the reference browser, Amaya and, with INRIA, ran ESPRIT projects to set up W3C national offices in Europe and to spread web awareness through the continent. ERCIM took over the running of the European Office of W3C in 2003 and Keith Jeffery, formerly of RAL, became President of ERCIM between 2007 and 2011.

To conclude, Bob showed a graph of the Atlas/RAL computing capacity from 1964 to 1999 in Ferranti Orion units, based as far as possible on statements of actual (rather than peak) performance that were made at the times when new machines or upgrades were installed. It showed an end-to-end increase of 777,056 (Bob noted the lack of an error bar!), and the growth trend was roughly consistent with Moore's Law.

50 years of the mathematical software library HSL

John Reid, of the Scientific computing Department's Numerical Analysis Group, described the scientific computing environment at Harwell in the early 1960s and the big improvements to the performance of scientists' own-written Fortran programme that were possible if they had access to a library of subroutines that incorporated good numerical analysis methods. This was the beginning of the HSL which grew to 90 subroutines, written in IBM Fortran, for dealing with approximation, ordinary differential equations, quadrature, linear algebra, sorting, optimisation and other techniques.

At first existing algorithms were coded, but the importance of associated research within the group was gradually realised and led to research of international standing becoming a key activity within the team. It covered optimization, approximation, stiff ordinary differential equations, and large sparse linear algebra, all with implementation as HSL subroutines. Demand for HSL to become available outside Harwell grew, as did the need for portability and the library was moved in stages to Fortran 77. In 1988 the matrix library was marketed by the Numerical algorithms Group (NAG) and there was a move from issuing single subroutines to packages of subroutines.

The group relocated from Harwell to RAL in 1990, switched to 2-yearly releases and from 1995 began to include Fortran 90 routines. Parallel programming using MPI was included from 2000 and using OpenMP from 2009. From 2011 Matlab and C interfaces were introduced for key routines. During this period free downloads became available to any academic for teaching and research purposes. There are now over 130 state of the art packages with a high standard of reliability and a strong international reputation. In the last three years there have been over 2000 downloads by more than 2000 academics from 68 countries, and the software is used in a huge number of applications.

The Present and Future

The Scientific Computing Department Today

David Corney, acting Director of the Department, which was formed in April 2014 and contained 180 staff supporting over 7500 users, summarised its main activities as: applications development and support; compute and data facilities and services; research, with over 100 publications annually; delivering over 3500 training days pa; systems administration, data services, high-performance computing; and, numerical analysis & software engineering. He outlined the large scale and the breadth of use of the computing facilities and gave examples of many of the specific activities.

His colleague **Barbara Montanari** described theory, experiment and simulation as the 'Three Pillars' approach to R&D, quoting several examples of disciplines where simulation had had major impact on understanding, and she outlined the growing dependence on a fourth pillar, data.

Derek Ross described SCARF (Scientific Computing Application Resource for Facilities), a powerful cluster system which was used by the Central Laser Facility, the ISIS neutron source at RAL and by external users as well as by the Scientific Computing Department itself, for development, simulation and analysis.

The four forces of change for intense computing

Cliff Brereton, Director of the Hartree Centre, discussed the beneficial economic impact that software and modelling could have in helping design better products faster and more cheaply. He described the Centre's very powerful computing facilities, the supporting infrastructure of expert staff and services, and the broad range of clients and collaborators.

The four forces of change driving intensive computing are industrial engagement, the computational power required, managing and providing analytics for big data, and democratisation. Major concerns for the near future were that Moore's Law appeared to be coming to an end and that present computing and data architectures are not sustainable into the long term. Emerging platforms and new architectures will need new algorithms, languages, compilers, middleware and applications. Analytics will realise the value of a data driven world.

Future Directions for Computing in Particle Physics

David Britton of the University of Glasgow and Project leader of Gridpp, which is responsible for the UK's computing for the Large Hadron Collider (LHC) at CERN, noted that at present the user community is dealing with over 200 Petabytes of data on disk and this will grow by more than a factor ten by 2027. The complexity of events in terms of the numbers of vertices will increase more than six-fold in the same period and, because the associated computational load increases superlinearly with that number, a growth rate in computing power is required that far exceeds the Moore's Law growth rate.

Innovative solutions for smarter working are being sought. The present hierarchy of Tier 0, Tier 1 and Tier 2 computing facilities will become more of a mesh architecture. 'Analysis trains' can provide a more effective method of coordinating many separate user analyses, thereby reducing 'chaotic' analysis. For monte carlo work a technique of fast simulation can be combined with full simulation in analysing events in an integrated framework that makes a hundred-fold speed up possible.

For computing hardware, particle physics uses commodity devices which are driven by external market forces. Memory sizes per core do not necessarily match particle physics requirements and work is being done to reduce memory demands by 'threading'. Memory savings can be huge but at the expense of more difficult programming and the need for conversion of legacy code; a further factor is that manpower is increasingly expensive compared with hardware. Other potential developments include the better joining up of various e-infrastructures and evolving the present particle physics grid to use Cloud technologies where appropriate.

Future Directions in Gravitational Physics

Stephen Fairhurst, of the School of Physics and Astronomy at the University of Cardiff, noted that the production of gravitational waves by accelerating masses was a feature of Einstein's general theory of relativity, and that experimental data had been gathered in the past decade that closely matched the theoretical predictions in the case of one particular rotating binary system. He outlined the different types of wave phenomena (continuous stochastic, short bursts) that different types of astronomical source could produce and the huge challenges in modelling them computationally.

As an example, the waveforms produced by collisions between black holes could be modelled analytically when the holes were far apart but, as they grew closer to merging, non-linear Einstein equations had to be solved numerically. The spin of black holes, which caused precession of the colliding system, added a complicating factor. A naïve coverage of parameter space would need about 10,000 simulations and, since at present one simulation takes one month on a 256-core computing facility, the computational requirements in this field were vast.

On the experimental side, he described the global network of detectors now searching for further evidence of gravitational waves and how future detectors are expected to have a ten-fold increase in sensitivity. The data will be full of noise transients that can mask signals. During 2015-16 there will be a need to search through months or years of data using about 10,000 template waveforms, increasing to millions of waveforms incorporating spin and precession effects by the end of the decade. So again the computational requirements are huge.

Future Directions In Digital Chemistry

The theme of this talk by Jeremy Frey, Professor of Chemistry at the University of Southampton, was how to reduce and manage uncertainty in the face of the vast quantities of data made available through modern Information technology. The kind of questions a researcher might have in mind include: what is already known about a topic?; who is doing what? how was a result obtained?; how to access the data?; and how has data been integrated? Since research Is increasingly inter-disciplinary, help is needed in information exchange, and today data is essentially lost if it is not available to a computer. In practice experiments are often repeated, data might be stored locally on a computer and unable to be found, handwriting might be unreadable, laboratory notebooks might be lost or damaged or the correct data might not have been recorded first time round.

Professor Frey described a number of projects that attempt to come to grips with these issues and noted that there is a blogosphere of scientific communication as well as conventional publication journals. He described Big Data in chemistry as being big, broad, bothersome and broken. Open

Notebook Science has obvious benefits in terms of access and collaboration but is not always the way to work because of IPR and related issues.

There was a pressing need to speed up the knowledge discovery process and to change the way in which experiments are designed and built. In times of change it is even more important to bring dissemination into the laboratory – it is an essential part of research.