Atlas 50th Anniversary, December 2012.

NRDC and the case for a British supercomputer, 1956 – 1960.

Simon Lavington, November 2012.
Comments welcomed to: lavis@essex.ac.uk

Acknowledgements.
The reminiscences of several former Atlas pioneers have contributed to the compiling of these notes. Particular thanks are due to Dai Edwards and Bob Hopgood.

Introduction.
The major players during the period 1956 – 1960 period were:

- the United Kingdom Atomic Energy Authority (especially Sir John Cockcroft at Harwell and Sir William Penney at the Atomic Weapons Research Establishment). Further details of the UKAEA organisation are given in Appendix 1.
- the National Research Development Corporation (and especially Lord Halsbury, the NRDC Director and Christopher Strachey, a key NRDC technical expert);
- the Royal Radar Establishment (and especially Sir Owen Wansbrough-Jones, Chief Scientist at the Ministry of Supply (MOS));
- Ferranti Ltd.;
- EMI Ltd.
- The University of Manchester's Department of Electrical Engineering, situated in a new (1954) building in Dover Street, Manchester.

In the background was the Brunt Committee of the Department of Scientific and Industrial Research (DSIR), set up in 1949 "to keep under review the progress in the construction and use of [computers], to examine the fields in which they are likely to prove most useful, and to recommend the most promising types for further development". The Brunt Committee, and indeed all the above organisations, sought advice from the leading computing experts in government and academic research establishments – particularly from the universities of Cambridge and Manchester.

The National Research Development Corporation (NRDC) had been set up in 1949. Prior to 1956 NRDC had been involved financially in encouraging the development and marketing of several British-designed computers, particularly the Ferranti Mark I*, the Elliott 401 and the Ferranti Pegasus. From 1957 onwards, NRDC began to involve itself in promoting the idea of a new high-performance computer which, because of its ambitious performance target, became known informally as a *Supercomputer*.

NRDC soon took the lead in attempting to organise a British Supercomputer project on its own terms. Events did not go smoothly for Lord Halsbury, NRDC's Director. Writing towards the end of 1958, he said: "Were we [the UK] strong enough to compete? Ought we to try? Could we afford not to? Could any such proposal be established on a commercial basis? During the last two years I have unsuccessfully wrestled with divided counsels on all these issues." [ref.1]. For Ferranti at least, NRDC's efforts were not wholly appreciated. At one stage Sebastian de Ferranti suggested that the British Supercomputer project should be called *BISON*: Built In Spite Of NRDC [ref. 2].

Here are the main events in the protracted saga that eventually led to the emergence of the Ferranti Atlas as the UK's first supercomputer. Because of the sequential nature of NRDC's discussions, the account is most easily represented as a time-line. The information is principally drawn from [refs 3 and 4], with additional sources noted at the end of this document under *References.*

Timeline.

1956 Feb Ferranti's computer department moves into new premises at West Gorton, Manchester. With Pegasus and Mercury machines being manufactured, West Gorton soon became the largest computer production plant in Europe.

1956 May: After a visit to the USA by A S Douglas (Cambridge), J H Wilkinson & D W Davies (NPL) and J Howlett (UKAEA) under the auspices of the Brunt Committee, the Committee was advised that the UK was lagging behind the USA in the provision of scientific computing facilities [Ref. 1]. A comparison of the speeds of various contemporary US and UK computers is given in Appendix 2 below.

1956: Ferranti fails to sell a Mercury computer to AWRE; AWRE buys an IBM 704 instead (see Appendix 1). This stimulated Ferranti to propose to AWRE the development of a computer considerably faster than Mercury at a lower price. Though AWRE expressed interest in such a development, AWRE was unwilling to award a contract for its realization. Ferranti then worked out a more concrete proposal, suggesting a machine at least ten times as powerful as Mercury, which would be made even faster by linking several together; this enhanced proposal was submitted to the Atomic Energy Authority in April 1957 (see below). This proposal envisaged a machine developed entirely by Ferranti, to be completed in 1961.

1956 autumn Two Ferranti engineers are seconded to Manchester University to join Tom Kilburn's computer design team in the Department of Electrical Engineering, University of Manchester. Ferranti regarded the secondment as a means of 'keeping in touch'. A few months later Kilburn's team begins the MUSE (*Microsecond Engine*) high-performance computer project. From the start, Kilburn envisaged that the University would need the cooperation of industry because of the size of the envisaged machine. He hoped that, after some progress with the design had been made in-house, industry and/or government would be convinced of its viability and would finance its completion. He had a fall-back position of a fast but smaller sub-set machine, to be built with £50K made available from the Department's own research fund. This fund had been accumulating since 1951, as an agreed percentage of the University's overall *Computer Earnings Fund.* The initial 'earnings' came from the hire of computing time on the University's Ferranti Mark I machine.

1956 Dec. The design objectives for the Univac LARC and the IBM STRETCH computers are presented at the Eastern Joint Computer Conference in the USA. (See also Appendix 2 for technical details).

1957 Jan. Jack Howlett (AERE Harwell) and J Corner (AWRE) produce for the UKAEA a report entitled *The case for a big computer*, stating that 'by 1960 the AEA would need a computer comparable to LARC and STRETCH'. They assumed that the development costs (estimated optimistically to be between £0.5m and £1m) could be born by the UKAEA. This Report was submitted to Sir John Cockcroft (Head of Harwell) and to Sir William Penney (Head of AWRE). Penney was sceptical. At the end of January Cockcroft (who was enthusiastic) sent the report to NRDC.

1957 Jan. In response to the Howlett/Corner report, the NRDC Computer Sub-Committee asks Halsbury to review and compare US and UK scientific computer development activity. From about this time, the word 'supercomputer' starts to be used by NRDC.

1957 Feb. E H Cooke-Yarborough and Patrick Blackett (both at Harwell) produce a counter-proposal to that of Howlett and Corner. Cooke-Yarborough and Blackett said that head-on competition with STRETCH was not a realistic option at this time. Instead, the UK should establish an 'Institute for research into electronics' and develop a new generation of components for computers beyond STRETCH.

1957 March Halsbury suggests to the NRDC Board that NRDC should spend £1m over four years on the development of a high-speed computer. The Board was sympathetic but only on the condition that the UKAEA agreed to buy the resulting machine.

1957 2nd April Halsbury tells Cockcroft of the Board's decision (see above), proposing that NRDC itself (meaning Christopher Strachey) would design the machine under an advisory panel composed of F C Williams (Manchester), M V Wilkes (Cambridge), A M Uttley (RRE) and E H Cooke-Yarborough (UKAEA). NRDC would then place development contracts with universities, government establishments and UK manufacturers. The aim was to 'produce a computer roughly comparable to STRETCH by 1961 (ie one year after STRETCH was planned to become available).

1957 30th April Cockcroft calls a meeting of UK computer users and designers (including Kilburn, Wilkes & Strachey) at UKAEA to define requirements for a high-speed computer. Attendees included representatives from the UKAEA (including Howlett, Brian Flowers & Cooke-Yarborough), NRDC, the universities of Cambridge and Manchester, the National Physical laboratory (NPL), etc. The medium-term Harwell requirements were for a fast store of 10^5 numbers and 10^5 mesh-operations per second. (The 'mesh' referred to three-dimensional field calculations, so that one elementary mesh-operation was roughly equivalent to about ten machine instructions, one of which was a multiply). At the 30th April meeting, UKAEA found itself in disagreement with NRDC's 2nd April proposal (see above) on three grounds: (a) they disliked the implication that the machine seemed to be heavily dependent upon Strachey's ideas for the logical structure and only secondarily dependent upon improved electronic components; (b) they disliked NRDC's proposed form of project management; (c) they were suspicious of NRDC's hope that it could cover its investment by the sale of one machine to AEA. They felt that UKAEA might as well sponsor the project directly and keep control over it. But see later (June) for Sir William Penney's negative response about UKAEA participation.

1957 April Ferranti submits a proposal (referred to above, under the entry for 1956) to AWRE. AWRE replied in July 1957, saying that it needed this new machine within 18 months and that, if later, Ferranti would have to compete with the UNIVAC LARC and/or the IBM STRETCH. AERE Harwell was also interested. In the event, NRDC took over from UKAEA the responsibility for organising the proposal for a new machine – (see below).

1957 spring to autumn Various discussions take place both internally within UKAEA, between UKAEA and NRDC, between UKAEA and MOS (principally RRE), and between NRDC and everyone else (including the Board of Trade). The likelihood of Harwell going it alone quickly faded. NRDC still wanted to take control of the project but few outside of NRDC were happy with this. Result: deadlock.

1957 June Whilst attending the Association for Computing Machinery (ACM) National meeting in the USA, Christopher Strachey 'personally visited both the LARC project at Univac in Philadelphia and the STRETCH project at IBM Poughkeepsie' but was unable to get firm information because both companies were secretive.

1957 June Sir William Penney writes to Halsbury, in effect saying that there was no prospect of UKAEA placing an advance order for a British supercomputer. UKAEA had two main branches (see also Appendix 1): basic research (at Harwell) and weapons (at AWRE Aldermaston). Penney said that AWRE would certainly need a fast computer but, with their military priorities, AWRE had to take the first to become available as a proven, reliable, product, be it British or American. The case for Harwell to have a fast computer had not yet been approved. In conclusion, Penney told Halsbury that the UKAEA should play no active role in the development of computers

1957 26th June NRDC agrees to put 'not less than £1 million' into a high-speed computer project. At this point, NRDC felt that the development would be beyond the capabilities of any one single UK manufacturer. Accordingly, NRDC proposed a joint project which would (a) have contributions from all the UK's centres of expertise, and (b) would give NRDC an opportunity to gain some return on its investments. In the event, it proved impossible for NRDC to satisfy (a) and (b) together. (Note that by this time, IBM was said to be spending $28m annually on STRETCH and to be deploying 300 graduates on the project).

1957 Sept. By now, NRDC was proposing to set up a subsidiary company with £1m capital, which would then subcontract out the necessary R&D and manufacturing of a supercomputer.

1957 Nov. NRDC and its advisors (which included F C Williams & Maurice Wilkes) disagree on the details of a subsidiary company.

1958 22 Jan Blackett put to NRDC his proposal that RRE run the project on behalf of MOS. RRE's conditions for doing this included the proviso that a single, well-established, company should be introduced as the prime contractor at an early stage of the project. Wansbrough-Jones, the head of RRE, also insisted on establishing a technical consensus before going ahead. To test the opinion of the leading experts, the following meeting was held:

1958 27th Feb. A meeting was held between NRDC and representatives from AERE, RRE, NPL, MOS and the Universities of Cambridge and Manchester. This meeting, later to be known as the *Second Harwell Computer Conference*, recommended the development of two computers: (a) a shorter-term project, having a fast store of 10^5 numbers and 10^5 mesh-operations per second; (b) a longer-term project involving new components and new techniques. At this point, the Manchester University MUSE project was considered a basis for (a). Project (b) was put on one side, to emerge much later in 1962 as the Advanced Computer Techniques project.

Meanwhile, a Working Party consisting of Kilburn, Wheeler, Strachey, Howlett and Cooke-Yarborough was set up to check the suitability of MUSE for (a), the short-term project. The Working Party concluded that MUSE was sound but its performance would fall short of Harwell's needs – especially in respect of storage capacity and peripherals. With a bit of re-designing and enlarging, it would be suitable. However, Kilburn refused to redesign MUSE. Also, on the matter of who should be in charge of the project, Kilburn felt that he, as the lead engineer, should take final responsibility. He was unwilling to be placed under the control of someone else from outside Manchester, (for example someone from NRDC or MOS). Kilburn contributed to the Working Party's discussions but, for obvious reasons, did not participate in the writing of the final Report. It was Strachey who drafted the final Report. Excerpts from its contents have been reproduced in [ref. 11].

1958 March Ferranti decide to restrict its own independent development of a large machine. Instead, Ferranti was to go for a medium, business data-processing, machine (Orion), and to cooperate with Manchester University on a larger (scientific) machine, later to become known as Atlas. (Much later in 1973 Peter Hall, the Manager of Ferranti's Computer Department from December 1958, said that "Orion was not so successful as it should have been because of lack of effort – which went into Atlas").

1958 June Cockcroft is anxious that the collaboration between Manchester University and Ferranti should be allowed to continue. NRDC did not really want this to happen, since the project was said to be incompatible with NRDC's long-term interests. However, they would have acquiesced if the project was put under the control of RRE Malvern as the design and contracting authority responsible to NRDC. Cockcroft then had talks with Williams and Ferranti. They hammered out a proposal which would have involved NRDC passing £1m to the project through RRE and then relinquishing direct control of what followed. RRE was to be the Design Authority, advised by a committee of independent experts. Manchester University would retain the prototype. The method whereby NRDC would recover its development money was left open. Most of the NRDC committee seemed to be inclined to favour this compromise but Halsbury did not like it and recommended deferring a decision till the autumn. Despite this:

1958 6th August: Sir William Black, Chairman of the NRDC Board, writes to Ferranti, asking if Ferranti would be interested in being the main contractor for a supercomputer project. Vincent de Ferranti informed Black that his company would "like to be made the main contractors for a project of the nature suggested as long as arrangements could be made which would not put us in the bankruptcy courts". Subsequent negotiation over the next few weeks touched on issues such as: how should the risks be divided? By what method, and over what time, should NRDC's loan be repaid? Who should be the 'design authority' for the project? Who would own the first production model to be built? No agreement was reached. NRDC then approached English Electric Ltd., to enquire whether English Electric would be the main contractor. This approach produced a negative answer. By the autumn of 1958, NRDC was close to giving up the idea that it could ever launch the supercomputer project.

1958 mid-October UKAEA informs NRDC that UKAEA could not now support a project to build a supercomputer in the UK and that, consequently, UKAEA felt free to consider purchasing such a machine from the USA. There were three reasons for this: (a) UKAEA realised that NRDC was getting nowhere with its negotiations with UK industry and UKAEA's computing requirements could wait no longer; (b) IBM had given UKAEA further details of STRETCH, which seemed to them to make the Manchester University/Ferranti project redundant; (c) Cockcroft, who strongly favoured MU/Ferranti, had just left UKAEA to become the first Master of Churchill College Cambridge.

Thus, by the autumn of 1958 three things were apparent: (i) there was dissention within NRDC about the form of project control and recovery of investment: (ii) the main customer (UKAEA) had dissociated itself from the project; and (iii) a contractor had not yet signed up. Nevertheless, NRDC decided to try one last time – possibly because NRDC could see that the Manchester University project was making progress and that Ferranti was interested in it. Also, with UKAEA out of the picture, Ferranti had more justification for wishing to take more part in design studies and assessing the needs of potential customers other than the UKAEA.

Over the period March 1958 to June 1959 the Manchester University project had indeed made progress and had become better defined – (see [ref. 11]. Thus it had become less reasonable for NRDC to demand to take technical control. 'Technical control by NRDC' meant, in practice, control by Christopher Strachey. Hendry says [ref. 4]: "Strachey was a brilliant logical designer but he was not an engineer, had very strong personal views on the question of computer design, and appeared distinctly intolerant of other peoples' views". Maurice Wilkes said of Strachey in 1983: "It was characteristic of him to insist on minor points of difference with the same force that he insisted on major points. This made conversation with him somewhat exhausting" [ref. 5]. Upon leaving NRDC in the spring of 1959 (see below), Strachey set up as a private consultant from 1st June 1959. NRDC, Ferranti and EMI all paid him a retainer.

1958 Oct. Ferranti decided to go ahead with a cooperative project with Manchester University, the project thereupon changing its name from MUSE to Atlas. This decision was confirmed in January 1959. (The slight uncertainty about informal versus formal agreement may have arisen because Brian Pollard resigned as Manager of Ferranti's Computer Department and Peter Hall took over in December 1958 – bringing much-needed dynamism). From this point on, Tom Kilburn and Peter Hall exercised joint control over the project – both in respect of development carried out at the University and at Ferranti's West Gorton facility. The terms of the agreement, which is set out in full in [ref. 11], were as follows: (i) the first machine was to be installed at Manchester University, who would have half the operating time; (ii) Ferranti would operate this machine at their own cost and would sell time to outside users in industry, commerce and academia; (iii)Ferranti would pay Manchester University £10K per annum plus a percentage of the income gained by selling time on the computer; (iv) the MUSE team of between 12 and 20 academics was increased in due course by the addition of between 30 and 40 Ferranti employees. The £10K went into the University's *Computer earnings fund* which had been started in 1951. At the end of the Atlas agreement in 1972, this fund stood at approximately £300K. In 1973 Peter Hall estimated that the total cost of the Atlas project had been £2m, of which Manchester University had originally put in £50K (plus academic staff wages and overheads).

1958 Oct. NRDC made fresh approaches to EMI, English Electric and Ferranti. English Electric and EMI were the first to respond, EMI positively and English Electric negatively – (Ferranti was late in responding, due to the Pollard/Hall changeover).

1958 Dec. NRDC decides to go with EMI, who were content with NRDC's grant-recovery proposals. Then in January 1959 Ferranti submitted its response to NRDC's financial proposals. NRDC deliberated about EMI and Ferranti for some weeks.

1959 Feb. Strachey files a patent for 'time-sharing in large computers' (British Patent 924672, granted in 1963). In 1959, 'time-sharing' was the term used in the UK for what we would now call 'multiprogramming'. Strachey's patent made 57 claims, specifying a director (supervisor) in fixed store, a memory-protection mechanism and an interrupt priority structure. In 1985 an expert said that "It is not currently possible to assess the significance of Strachey's patent in terms of either priority or commercial use" [ref. 5]. Note that Kilburn's group at Manchester University filed 15 Atlas patents between October 1957 and February 1962 through NRDC as their patenting agent.

1959 17th March This NRDC meeting was Halsbury's last, since he was soon to retire. Strachey left NRDC at the same time. At this 17th March meeting, the NRDC Committee was asked to decide between two schemes for a supercomputer:

a) a Ferranti scheme: £510K grant; Kilburn to make technical decisions; repayment to be as a percentage of company's profits; project duration = 3 to 4 years.
b) an EMI scheme: £280K grant; Strachey to advise how the EMIDEC 2400 computer could be improved; repayment terms to be as levy on *all of* the company's computer sales; project duration = 2.75 years.

Whereas Strachey had previously said that the Manchester University/Ferranti proposal was ahead of EMI's offering, he now said that the two current proposals had equal technical merit. (The background to EMI's efforts, the EMIDEC 2400 and Strachey's involvement with EMI, are given in Appendix 3 below). On 17th March the NRDC Committee favoured EMI 'by the narrowest of margins'. After the Committee meeting, Dennis Hennessey (NRDC's Deputy Director) expressed concern to Sir William Black (Chairman) about Halsbury's bias. As a result, the NRDC Board decided in the end to compromise and back both horses but with lesser funds. £250K was to be loaned to EMI and £300K to Ferranti, both loans to be recovered by a levy on the total computer sales (not profits) of each firm. The EMI supercomputer project effectively came to nought (see Appendix 3). Ferranti was already committed to the Manchester University project before NRDC offered the grant, and would have gone ahead anyway.

1959 May NRDC loaned Ferranti £300K towards cost of Atlas development, to be recovered from sales. Ferranti had paid off their loan by March 1963, according to [ref. 6].

1959 15th/20th June: IFIP UNESCO Computer Conference in Paris: separate papers on high-performance computers are presented by Strachey [ref. 10b] and by

Kilburn [ref. 10a]. Strachey's paper came as a complete surprise to Kilburn and the MUSE team when they arrived in Paris. The American company Control Data Corp. (CDC) subsequently told Ferranti that Kilburn's Paris paper "had saved them a great deal of time and money because it told them of avenues which had been fruitlessly explored and stones which were not worth turning over" [Ref. 2].

1959 autumn Ferranti and ICT have talks (unsuccessfully) about a merger of their computer businesses. Main obstacle: the clash of cultures between the technology-led Ferranti camp and the market-led ICT camp. Merger talks were renewed in 1961, and then finally concluded in 1963 [ref. 6], at which date ICT absorbed Ferranti's mainframe interests (including Atlas) but left Ferranti's smaller machine developments in the Process Control and Defence sectors.

1960 30th Sept. By this date, the outline design of all nine Atlas sub-units had been completed and the detailed logical design of 80% of these units had been finished. The package-count for 60% of the computer had been completed and the backwiring schedules for 20% of the computer had been produced. For further details, see [ref. 11].

1960 Nov. Three Atlas hardware papers (storage and ALU technologies) were published in the IEE Proceedings. These papers were presented by the Manchester University team at the IEE meeting on 1st March 1960.

Postscript.
The first production Atlas was ceremonially switched on by Sir John Cockcroft at the University of Manchester on 7th December 1962. At the time, it was rated as the most powerful computer in the world. On the 22nd September 1961 Ferranti hosted a dinner at the Savoy to celebrate UKAEA's order to purchase an Atlas computer for about £3.1m. This machine, the largest of the Atlas installations, was delivered to the National Institute for Research into Nuclear Science at Harwell, Chilton, Berkshire in nineteen trucks in April 1964, after a custom-designed building had been completed. A third Atlas was installed at the University of London in 1963. Meanwhile, Ferranti and the University of Cambridge collaborated on the design of a smaller version of Atlas, to be called Atlas 2. The prototype Atlas 2 was installed at Cambridge, where it came into operation in 1964. It was known as Titan. Two more Atlas 2 machines were delivered: to AWRE Aldermaston in 1964 and to the Ministry of Technology's Computer-Aided Design Centre at Cambridge in 1967.

Appendix 1.
The organisation and computational facilities of the UK Atomic Energy Authority.

In the period 1954 – 1974 the United Kingdom Atomic Energy Authority was a crucial player in setting the standards for Britain's high-performance scientific computing facilities. These standards were in turn driven by the complexity and strategic importance of reactor calculations. In the years relevant to the birth of Atlas (1956 – 1960) the phrase *Atomic Energy Authority* was used informally as an umbrella term covering the following principle organisations, ordered according to their computing needs:

(a) the Atomic Weapons Research Establishment at Aldermaston, founded in 1950;
(b) the nuclear reactor and weapons construction facility at Risley;
(c) the Atomic Energy Research Establishment at Harwell, founded in 1946;
(d) the reactor R&D centre at Winfrith, founded in 1958;
(e) the Centre for Fusion Energy at Culham, founded in 1960.

In the run-up to Atlas there was a distinct split between the immediate Cold War computing requirements of AWRE and the longer-term requirements of the other four organisations. AWRE was the front-runner. In 1954 AWRE obtained a Ferranti Mark I*, which was one of the first production computers available anywhere. Then AWRE acquired (or rented?) an IBM 704 in 1957 and an IBM 7090 in 1960, because of the need to process bomb calculations with some urgency. Meanwhile, AWRE was also pushing for an IBM 7030 (STRETCH) computer, which was finally delivered (for rental) in 1962 but with relatively little accompanying software. AWRE's cast-off facilities, particularly the IBM 704, were passed in due course to Risley, so the Risley computing provisions were usually relatively well-covered. Ferranti Mercury computers were installed at Harwell and Risley in 1958. During the Atlas project run-up, Harwell also obtained modest amounts of computing time from both AWRE and Risley.

Appendix 2.
A comparison of the performance of representative US and UK production computers suitable for scientific applications, 1956 – 1962.

A2.1. Instruction times.

Instruction	*Ferranti Mercury (1957)*	*LEO III (1962)*	*EMIDEC 2400 (1962)*	*Elliott 503 (1963)*	*English Elec. KDF9 (1963)*	*Ferranti ATLAS 1, (1962)*
FXPT ADD	60	34 - 49	24 - 31	7.2	1	1.59
FLPT ADD	180	-	-	13.2	6 - 10	1.61 – 2.61
FLPT MPY	300 (360)	-	-	?	14 - 18	4.97
FLPT DIV	?	-	-	?	30 - 35	10.66 – 29.8

Table A2.1: Instruction times, in microseconds, for various British computers that were first installed between 1957 and 1963.
It is believed that neither LEO III nor the EMIDEC 2400 had floating-point hardware. The Ferranti Mercury was one fifth the cost of an IBM 704 computer and about one third of the speed.

Univac LARC, (1960)

Instruction	*IBM 704, (1955)*	*IBM 7090, (1959)*	*Univac LARC, (1960)*	*IBM Stretch, (1961)*	*Ferranti ATLAS 1, (1962)*	*CDC 6600 (1964)*
FXPT ADD	24	4.8	(4 ?)	1.5	1.59	?
FLPT ADD	84	16.8	4	1.38 – 1.5	1.61 – 2.61	0.3
FLPT MPY	204	16.8 – 40.8	8	2.48 – 2.7	4.97	1
FLPT DIV	216	43.2	28	9.0 – 9.9	10.66 – 29.8	3.4

Table A2.2: Instruction times in microseconds for various American computers that were first installed between 1954 and 1964, compared with Atlas.

A2.2. More details of STRETCH (also known as the IBM 7030).

IBM's goal was to produce a performance two orders of magnitude faster than the IBM 704 computer. The company's initial conceptual thinking for STRETCH, for example the instruction look-ahead strategy, dates from 1955. The project was formally initiated in January 1956 and the contract with the Los Alamos Atomic Energy Commission Laboratory was formally signed in November 1956. The first STRETCH was delivered to Los Alamos in April 1961, accepted in May 1961, and used until June 1971. Seven more were delivered, including to the Lawrence Livermore Laboratory (1961) and to AWRE in England (1962). STRETCH was initially priced at $13.5 million, finally reduced to $7.78 million each.

The STRETCH architecture was characterised by significant amounts of overlapping, low-level concurrency and multiplexing, with the necessary instruction buffers and operand buffers. The instruction look-ahead mechanism fetched a group of instructions and partially interpreted them – work that might have to be undone later if the program thread was not straightforward. Without all these tricks, the basic memory performance was only expected to be 6 times that of the IBM 704 and the basic circuits about 10 times that of the 704.

Primary memory consisted of up to six separate core stacks, each 16K words, each 2 microsecond cycle-time. Typically, two stacks were 2-way interleaved and were used for instructions; four stacks were 4-way interleaved and used for data. The word-length was 64 bits. Instructions were either half-word (32 bits) or full-word (64-bits), single-address format, with 18 bits specifying the word address (thus enabling 2^{18} or 256K words to be directly addressed). There were 15 index registers. There was a rich instruction repertoire, of which some orders were quite complex, for example for array-processing. Decimal arithmetic was included (as well as the usual binary). 'Multiprogramming' was aided by a pair of datum/limit registers. IBM's STRETCH had a bigger, faster disk than the UNIVAC LARC moving-head drum, but could achieve only a few accesses per second, having only a single access mechanism.

In their seminal book on computer architecture, Bell and Newell comment thus: "STRETCH was not outstandingly successful – only a few were built at a cost substantially exceeding their contract price and with a performance only modestly better than the art at the time of their production" [ref. 8]. Bob Hopgood, who wrote compilers for both STRETCH and Atlas and implemented a large Quantum Chemistry package (MIDIAT) on both computers, has said [ref.9]:

"STRETCH could run extremely fast if you had the code set up just right and it remained in core memory. It had some terrible deficiencies as well. It made guesses as to which way a conditional jump would go and if you got it wrong it had to backup all the computation it had done. So the same conditional jump could be as much as a factor of 16 different in time between guessing right and wrong. The STRETCH nuclear weapon codes at AWRE Aldermaston probably outperformed Atlas by quite a bit. On the other hand Atlas ran some large number theory and matrix manipulation calculations much faster than STRETCH. 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'".

A2.3. More details of the Univac LARC.
'LARC' originally stood for *Livermore Automatic Reaction Calculator***,** the title of a proposal sent out for comment to IBM and to Univac by the Lawrence Radiation Laboratory at Livermore. Subsequently Univac changed LARC to stand for *Livermore Automatic Research Computer***.** Only two Univac LARCs were built, the first of which was delivered to the Lawrence Livermore Laboratory in March 1960, though it was not until March 1961 before it passed its acceptance tests. The second LARC went to the US Navy.

LARC had decimal arithmetic (4 bits per decimal digit), with a 48-bit word length. The basic configuration had one CPU but a second could (in theory) be added. The primary memory (core) had eight banks, each of 2500 words, giving 20,000 words in all. This was expandable to 97,500 words. The access-time was four microseconds. There was provision for movable-head drums (between 12 and 24) and tape decks (4 to 40). Each of 12 drums in the first LARC installation held 250,000 words. LARC was built using surface barrier transistors, which were becoming obsolete by the time the first LARC was delivered.

Appendix 3. The EMIDEC 2400 and 3400 and the involvement of Christopher Strachey.

The following notes are largely taken from John Hendry's book ref [4].

Electric and Musical Industries Ltd. (EMI) decided, with NRDC encouragement, to enter the computer business in 1954/55. After building a one-off machine for the British Motor Corporation, things gathered momentum. Active work on EMI's EMIDEC 1100 machine started in 1957 and about 25 were sold – see:
https://www.ourcomputerheritage.org/Maincomp/Emi/ccs-m1x1.pdf
The relationship between EMI and NRDC soon became particularly good because:
(a) EMI was new to the computer game and was willing to accept NRDC's (and especially Halsbury's) strategic guidance on the national needs. In 1955 Halsbury was keen to provide a UK competitor for the IBM 704 and 705 computers and their successors.
(b) Consequently, EMI was also willing to accept NRDC's technical guidance. In the summer of 1955 Strachey helped EMI work out detailed proposals for a new computer that was to become the EMIDEC 2400. By 1957 this project had got well under way.
(c) Unlike Ferranti, EMI was willing to accept NRDC's payback scheme of a small levy on *all* EMI's computer products over a long period. The final EMIDEC 2400 contract with NRDC was signed in October 1956, with NRDC putting up three quarters of the estimated development costs of £320K.

As has been seen, NRDC started its quest for a British supercomputer project early in 1957. In this context it was not until the autumn of 1958 that NRDC began to look towards

EMI, whose EMIDEC 2400 project was by then overrunning both on cost and time. In October 1958, following an enquiry from NRDC, EMI replied that it would be willing to undertake a large fast computer project. In December 1958 EMI's EMIDEC 2400 team, with considerable technical input from Chris Strachey (NRDC), made a new technical proposal to NRDC which became known as the EMIDEC 3400 project.

Strachey had expressed a dislike for the Manchester Atlas design and was particularly critical of the Atlas paging system, so he was clearly interested in alternative schemes. The 3400 project became Strachey's answer to Atlas. When advising the NRDC Committee, Strachey estimated that the 3400 computer would be:

(a) three times faster than the IBM 704 but significantly cheaper;
(b) about 0.25 – 0.5 as fast as Atlas;
(c) ready and working sooner than Atlas;

This EMIDEC 3400 project was essentially the EMI high-performance computer proposal that was backed by the NRDC Committee in March 1959.

The EMIDEC 3400 turned out to be a fiasco. By November 1959 EMI was proposing that the full design be postponed until the 2400 was completed, for which NRDC was asked for more money. The 3400 project continued to potter along. EMI sold its computer interests to ICT in 1962, after which the 3400 project lost momentum. The project was dropped by ICT in 1963, after the Ferranti merger with ICT.

It is relevant to explore Strachey's ideas for a supercomputer in a little more detail. One approach to this exploration is through an examination of Strachey's paper published at the UNESCO Paris Conference in June 1959 [ref. 10a]. With the benefit of hindsight, it appears that some of the ideas in Strachey's paper could be said to have been inspired by – or shared by – the MUSE/Atlas project. These include: the precise speed and use of a Fixed Store; the use of private RAM for the 'Director' working area; the arithmetic instruction speeds; the use of fixed-block tapes; the need for some kind of job control language. Two of Strachey's ideas that were definitely not in MUSE/Atlas include: the provision of three consoles (engineers, program testing, main operators); the use of datum/limit registers for memory management.

After resigning from NRDC in the spring of 1959 Strachey set up as a private consultant. Then in June 1962 he moved to the Mathematical Laboratory at the University of Cambridge to work for Maurice Wilkes full time on a new programming language and compiler for Titan. At the Mathematical Laboratory he joined David Barron and David Hartley, to work on their proposal for CPL (originally standing for *Cambridge Programming Language* but later becoming *Combined Programming Language* when collaboration with London Atlas arose). A slow and preliminary version of CPL was eventually introduced in 1966 but it does not seem to have had much production usage.

Strachey resigned his post at Cambridge on 31st July 1965 and also wound up his private consultancy. After a period as a visiting lecturer at MIT he moved to the Program Research Group at Oxford in April 1966. He died in 1975, aged 59. He is perhaps best known for his seminal work on denotational semantics – see [ref. 5].

References.

1. The Earl of Halsbury, *Ten years of computer development.* Computer Journal, vol. 1, issue 4, January 1959, pages 153 – 159.

2. B B Swann, *A history of the Ferranti Computer Department.* Privately 98-page typed document, privately circulated. (See National Archive for the History of Computing (NAHC) catalogue reference FER/C30).

3. P Drath, *The relationship between science and technology: university research and the computer industry 1945 – 1962.* Ph.D. thesis, University of Manchester, October 1973. (In addition to studying all relevant NRDC documents, Drath was able to conduct personal interviews with many of the key players).

4. John Hendry, *Innovating for Failure: government policy and the early British computer industry.* MIT Press, 1990. ISBN: 0-262-08187-3. (Hendry was able to study relevant UKAEA papers that were unavailable to Drath).

5. M Campbell-Kelly, *Christopher Strachey, 1916 – 1975: a biographical note.* Annals of the History of Computing, vol. 7 no. 1 January 1985, pages 19 – 42.

6. J F Wilson, *Ferranti: a history. Building a family business, 1882 – 1975.* Carnegie Publishing Ltd., Lancaster, 2001. ISBN: 1-85936-080-7.
Volume 2: from family firm to multinational company, 1975 - 1987, published in 2007 by Crucible Books, Lancaster. ISBN: 978-1-905472-01-7.

7. (a). S W Dunwell, *Design objectives for the IBM STRETCH computer.* Proc. EJCC, pages 20 – 22, 1956.
(b). E Bloch, *The engineering design of the Stretch computer.* EJCC, pages 48 – 59, 1959.

8. C G Bell and A Newell, *Computer structures: readings and examples.* McGraw-Hill, 1971.

9. F R A Hopgood, e-mail to Simon Lavington, 29th February 2012.

10 (a). C Strachey, *Time sharing in large fast computers*. Proceedings of the International Conference on Information Processing, UNESCO, Paris, 15th – 20th June 1959, pages 336 - 341.
(b). T Kilburn, *MUSE.* Proceedings of the International Conference on Information Processing, UNESCO, Paris, 15th – 20th June 1959, page 433.

11. S H Lavington, *Timeline of the MUSE/Atlas project at Manchester University, 1955 – 1971.* Document produced in November 2012, in preparation for the 50th Anniversary of Atlas. See:
http://curation.cs.manchester.ac.uk/atlas/docs/MUSE%20Atlas%20Manchester%20final.pdf

Simon Lavington lavis@essex.ac.uk *12/11/2012.*