

RESURRECTION

The Bulletin of the Computer Conservation Society

Computer Conservation Society

Aims and objectives

The Computer Conservation Society (CCS) is a co-operative venture between the British Computer Society (BCS), the Science Museum of London and the Museum of Science and Industry (MOSI) in Manchester.

The CCS was constituted in September 1989 as a Specialist Group of the British Computer Society. It is thus covered by the Royal Charter and charitable status of the BCS.

The aims of the CCS are:

- ◇ To promote the conservation of historic computers and to identify existing computers which may need to be archived in the future,
- ◇ To develop awareness of the importance of historic computers,
- ◇ To develop expertise in the conservation and restoration of historic computers,
- ◇ To represent the interests of Computer Conservation Society members with other bodies,
- ◇ To promote the study of historic computers, their use and the history of the computer industry,
- ◇ To publish information of relevance to these objectives for the information of Computer Conservation Society members and the wider public.

Membership is open to anyone interested in computer conservation and the history of computing.

The CCS is funded and supported by voluntary subscriptions from members, a grant from the BCS, fees from corporate membership, donations, and by the free use of the facilities of both museums. Some charges may be made for publications and attendance at seminars and conferences.

There are a number of active Projects on specific computer restorations and early computer technologies and software. Younger people are especially encouraged to take part in order to achieve skills transfer.

Resurrection

The Bulletin of the Computer Conservation Society

ISSN 0958-7403

Number 49

Winter 2009/10

Contents

Editor's Remarks	2
<i>Dik Leatherdale</i>	
Society Activity	3
News Round-Up	8
Pioneer Profiles – Douglas Hartree	10
<i>Crispin Rope</i>	
KDF9 Time Sharing: Eldon 2 is not EGDON!	15
<i>David Holdsworth</i>	
Elliott Computers for Motorway Signalling	24
<i>Richard Burwood</i>	
That Atlas Switch Panel	29
<i>Dik Leatherdale</i>	
System 25 – ICL's Cinderella Success of the Eighties	32
<i>Michael Knight</i>	
Forthcoming Events	39

Editor's Remarks

Dik Leatherdale

Ten years ago, the IT industry underwent something of a hiatus with the so-called “millennium bug”. On the whole, we acquitted ourselves rather well. The public utilities were uninterrupted. No aircraft fell from the sky. The merchants of doom were left with an embarrassing amount of canned food.

To this day, there are some outside the industry who still believe that the whole thing was a hoax and there wasn't a problem in the first place. Nick Davies, in his book *Flat Earth News*, devotes a whole chapter to the matter. But those of us who spend months and years grafting away, know better. People tasked with sorting things out told horror stories of systems which would, had it not been for our intervention, have failed spectacularly. At least one would have reverted to working in pounds, shillings and pence.

The assertion I heard at a Tandem User Group Y2K meeting that the millennium bug was the most serious crisis ever to hit the industry seemed, to me at least, to lack a sense of historical perspective. “Where were you on 15 February 1971?” I mused. But the answer was soon apparent when I looked at the faces around the room.

Yes, vast amounts of money were wasted “to be on the safe side”. I wasted a tidy sum myself. But, strangely, it's a tribute to what we achieved that some otherwise serious commentators still don't believe us.

2010, of course, is also the 50th anniversary of the Algol 60 Report, perhaps the most influential software document of all time. In case you thought we'd forgotten, the 50th edition of *Resurrection* will, with luck, come to your aid.

This edition's pioneer, Douglas Hartree, is now mainly remembered for his remark that the UK could get by with four computers. But there is much more to the man than that. Crispin Rope's magisterial rehabilitation of Hartree ranges widely across the achievements of this polymath (and great-grandson of Samuel Smiles) who passed away more than half a century ago.

Finally, the adoption of a new constitution brings with it a new statement of Aims and Objectives. Look inside the front cover.

Society Activity

Revised CCS Constitution – *David Hartley*

At an Extraordinary General Meeting of the CCS held on 15 October 2009, members of the CCS adopted a new constitution for the Society. Our parent body, the BCS, had changed its fiscal year, and the Society was therefore required to change the timing of its Annual General Meetings. In revising the constitution, which was last changed 10 years ago, the CCS Committee agreed to propose an updating to bring it in line with current practices and to rectify certain omissions. For example, the 1999 constitution made no mention of the Society's North West Group.

The BCS had also laid down certain new requirements on its Specialist Groups to ensure that they operate within the terms of the BCS's Royal Charter and to ensure that the charitable status that it enjoys is not jeopardised. One requirement is for members of BCS Specialist Groups to be members of the BCS, but it had been agreed that this could not apply to the CCS in view of its wide and open membership. However, the new CCS constitution requires that the Chairman, Secretary and Treasurer are all chartered members of the BCS.

The new constitution maintains a large committee with all leaders of Projects (formerly Working Groups) being members of the Committee *ex officio*, and with provision for the Committee to co-opt additional members as it thinks fit.

The new constitution is available on the Society's website at www.computerconservationsociety.org/Constitution%202009%20approved.pdf.

Society Committee

The eagle-eyed amongst our readers will have noticed that the list of committee members in *Resurrection 48* included two newcomers - Catherine Rushmore who now represents MOSI – the Museum of Science and Industry in Manchester - and Tony Frazer of the Harwell Dekatron Computer Project. In this edition Peter Onion also joins the list as the new leader of the Elliott 803 Project. All are most welcome additions.

Peter's appointment is a consequence of John Sinclair's decision to step down from the role. Kevin Murrell writes –

John Sinclair has led the Elliott 803 working party since its inception in 1989. He attended the inaugural meeting of the CCS and volunteered to lead the conservation project for the Science Museum's 803. The machine was moved to the Old Canteen in South Kensington and work

began in earnest, returning the system to working order. John later helped to move the machine to Blythe House. Since 1994 John has maintained another 803 computer at The National Museum of Computing. This machine now runs very reliably and is demonstrated twice a week. John is undoubtedly the expert in 803 engineering – including being able to detect failed transistors from the smell alone!

Harwell Dekatron Computer (WITCH) – *Tony Frazer et al*

Since the arrival of the Harwell Dekatron Computer at the National Museum of Computing in September this year, we have made considerable progress in the assessment of its condition and made a start on the restoration to full working order. The completeness and general condition of the machine surpassed our expectations, including all peripheral “organs”, cables, spare valves, schematic notes and punched tapes.

We decided that the power supply rack was a logical first step in the restoration.

First we examined the two GPO 110v DC supply units (for the reperforator and page printer). After removal of dust, replacement of perished rubber cables and safety tests, we successfully powered up both supplies (they are in parallel) with a Variac and with an Anglepoise lamp as a load.

Next we examined the STC 50v DC relay supply and found four capacitors dated 1947! We decided to leave these in position and replace them in circuit with a single modern device mounted adjacent to the originals. This unit has a rheostat connected across the output as a crude regulator to minimise voltage fluctuations due to external loading. After replacing the defective toggle switch and a fuse holder with identical units, we powered up the unit and noted that it was drawing about 1 Amp without an external load. After a short period of celebration and photographs, the main supply fuse blew! We will revisit this supply in the near future.

The HT Rectifier Unit presents a greater challenge. A preliminary test of the components indicated that the choke in the +370v circuit was open circuit. This was no great surprise, as the choke had evidently reached a high temperature at some point and the pitch with which it was potted had oozed from the choke casing, coating the interior of the chassis with a sticky brown residue. At this point, it seems evident that the HT supply was designed to supply sufficient current for four store units, but when an additional five store units were added, this caused

overheating of the HT Rectifier Unit over a long period of time with eventual failure.

Fortunately, we were able to find an almost identical unit, differing only in that it was mounted in side cheeks rather than potted in pitch. We removed the defective choke from its casing and replaced it using pitch salvaged from the original. The remainder of the HT Rectifier Unit was disassembled and cleaned, and we are in the process of reassembling it using the original components, new grommets and suitably rated plate wiring wire.

Another area of the project involves the recovery of software from an assortment of 5-hole punched tapes. We have read 90 percent of these using a slow optical tape reader over RS232 to a simple Visual Basic program on a PC. Examples of programs include various store clearing routines, arithmetic unit tests, Sine x, Fibonacci series, square and cube roots. We attributed difficulties in our initial attempts to recover the tapes to translucency of the paper; however it appears that the problem was more likely due to a build up of dust on the optical reader mechanism as frequently cleaning the unit resulted in much improved performance.

Looking to the near future, we hope to complete restoration of the power supply rack to working order early next year.

We will also be decorating the remainder of the hall in which the machine is now on display during restoration. We will install a low-level barrier with information panels, spotlighting and wall displays drawn from the vast amount of supporting material we have gathered so far.

Elliott 803 Project – *Peter Onion*

It has been some time since the last update on the work of the 803 Project so there is quite a bit of news. Over this summer the machine has been running well, but there have been some new faults to fix, and a couple of long standing ones have been tackled.

One of the large smoothing capacitors in the battery charger eventually revealed itself as the source of the heat that had been producing an occasional "hot smell". The decision was made to replace all six of these "beer can" sized capacitors. We found a supply of suitable replacement capacitors with manufacture dates in 1978 rather than 1963. Six of these were carefully reformed and John Sinclair installed them and tested out the supply.

The paper tape reader had started to take longer to "warm up" each day, but after a thorough service it now works from switch on. Some of the tape guide pins needed to be replaced as over the years the tape had worn notches into them and consequently they were no longer holding the tape in a steady position.

An OA6 diode on a logic board failed which put the accumulator arithmetic circuits permanently into subtract mode. Luckily the symptoms made this a fairly simple fault to track down.

The second core store had been disabled for a long time after a hard to find fault developed when the machine was decommissioned for an aborted move to another building. With the aid of a specially built signal probe the fault was eventually tracked down to a joint in the back wiring that had never been soldered!

With all 8K words of the store working again it has been possible to start running the Algol 60 compiler once more.

To punch a new copy of the two compiler tapes uses nearly a whole reel of paper tape. So to minimise the length that will need to be re-punched when a tape is ripped, the compiler has been split into seven shorter tapes.

Bombe Rebuild Project - *John Harper*

In *Resurrection 48*, I reported that the final components to complete the extra German Surface Fleet Enigma drums were available. The assembly of the remaining drums was completed in record time thanks to considerable team effort. All 39 were unit tested soon after. We then created a special test menu and were able to system test the whole batch in anger. We found just two errors. One was a misalignment problem on the internal wiring and was quickly fixed. The other was a loose connection that had not shown on unit test but did when the drum was rotated on the Bombe itself. Once diagnosed, this was very quickly fixed. We now have a full working set that enables us now to rerun Surface Fleet menus.

This now completes any planned construction. There are currently no plans to make any further parts or features but that is not to say that a few ideas that have come up over a period of time might not be progressed to the planning stage.

I would like to digress here a little and report how well the 3D small theatre in our area is performing. It shows a "how we made it" 3D slide show lasting about 12 minutes. It was developed by one of our team

members Mike Hillyard with help from his son. It is fully automatic even to the extent of turning itself on in the morning and off at night. Bletchley Park staff do not have to do anything except replace broken 3D spectacles from time to time. The whole system has turned out to be very reliable and all praise should go to Mike and his son.

In the last *Resurrection* I said that we now need help with providing more frequent demonstrations. We have had a healthy initial response but it is judged that we will need around six fully trained operators if Bletchley Park is to run regular public demonstrations, particularly at weekends. We have therefore produced a more comprehensive advert and this appears below.

Our website is still at www.jharper.demon.co.uk/bombe1.htm

The Bombe Rebuild Team now has a fully working and reliable machine but does not have enough resources to demonstrate it on a regular basis. We are now looking for volunteers to operate and demonstrate after appropriate training has been given. Demonstrations are most needed at weekends.

Many CCS people will have the appropriate skills and experience for this with further training. Training will be arranged at times when the Rebuild Team is on site at Bletchley Park.

The initial training will be Health and Safety orientated and candidates with a background in operating machinery and experienced in working with high voltages will have an initial advantage. H & S applies not just to those being trained but also to the precautions and procedures that need to be applied to safeguard the visiting public.

Following the H&S training a candidate will be shown how to interpret a menu, mount drums and plug up the back of the machine, etc. Jobs can then be demonstrated but it will also be necessary to identify the sort of basic problem that might be encountered. This is a similar level of competence that the WRNS operators achieved in WWII.

There is also a need in the longer term for volunteers to move from operator to maintenance mechanic where the level of knowledge will be similar to the RAF gentlemen who carried out this task at the outstations during WWII. Additional training will be made available for those wishing to progress beyond the operator level.

If you are able to help please contact John Harper – details inside back cover.

News Round-Up

Our congratulations to Steve Davies, the director of MOSI – the Museum of Science and Industry in Manchester. Steve, who recently gave a presentation to the Society’s North West Group has been appointed to the prestigious post of Director of the National Railway Museum at York and Shildon. We wish him well.

Tilly Blyth of the Science Museum has also just started a part-time role at the British Library as Senior Academic Consultant on their Oral History of Science project. The project will create a major archive in British science, including interviews from significant figures in the history of computing. She will combine this role with her current position as Curator of Computing and Information at the Science Museum. Tilly has been a good friend to the Society, so we wish her a rewarding experience.

101010101

On 30th September, Bonhams held an auction of a collection of early technology artefacts. Amongst the collection was a small group of pieces of the first LEO II to be delivered to an external customer, together with some associated documentation.



LEO II/3 was described by the vendor as “The first commercially sold computer”. *Resurrection* makes no comment! The lot was sold for £8,400 against an estimate of £2-3,000. Time to turn out the attic?

101010101

In *Resurrection* 47 we drew reader’s attention to the forthcoming IFIP conference at Brisbane. A call for papers on the subject of Computer History has now been issued and can be found at www.wcc2010.com/migrated/HC/HC_cfp.html.

101010101

There will be a reunion of LEO people on Sunday, 18th April 2010 in London. The date coincides with two significant anniversaries. It will be 60 years since Lyons applied for patents for the LEO project. The first programming team was also established in 1950.

The reunion is being organised by the LEO Computers Society. It will be held at The Park Inn in Southampton Row, London, starting at 11:30.

The event will mainly enable anyone associated with the supply and use of LEO Computers to meet with colleagues and others, to reminisce about their times working with computers, especially LEO, the earliest of which was the world's first business computer.

As usual there will be memorabilia on display and LEO videos will be shown. There will also be a short talk covering early computers and the work being done to preserve the heritage.

The aim is to keep the price below £20 which will include a cold and hot finger buffet lunch with coffee and soft drinks. There will be a cash bar.

Full details of the reunion will be circulated in the new year to all members of the LEO Computers Society. Attendance must be pre-booked through the Society. Anyone who was associated with LEO is welcome to join us. Please complete your details on the LEO website form (www.leo-computers.org.uk/newleoform.htm) or send a brief summary of your connection with LEO to John Hall (reunion@leosoc.freenetname.co.uk) and your name will be added to the mailing list. Those not online can contact John on 01480 474575. Please leave a message if there is no reply.

101010101

Sad to report the deaths of two Cambridge computing figures, both of whom took chairs at London University - Peter Landin at Queen Mary College and John Buxton at King's. Peter Landin was best known for his contributions to the theory of programming languages and their relationship to Alonzo Church's lambda calculus. John Buxton's varied career included spells at IBM and Warwick University. He (with the late John Laski) devised Control and Simulation Language (CSL) for Esso and was a contributor to the joint Cambridge/London CPL project which led in due course to the emergence of C.

Pioneer Profile: Douglas Hartree

Crispin Rope

Douglas Rayner Hartree (1897-1958) is unusual in that he had four claims to fame. Firstly, for developing a method for using quantum mechanics for detailed calculations at the atomic level (Hartree – Fock calculations – still very widely used). Secondly, Hartree worked in a number of other important areas of science including the cavity magnetron. Thirdly, he was a world leader – perhaps the foremost – in seeing the importance of numerical analysis at a time when this was considered rather a third class subject. And fourthly, Hartree was involved in the early days with five of the first very early computers.



Although Hartree's achievements in physics and numerical analysis are famous amongst those working in those fields, it is his early work in developing the new field of computing which concerns us here. Starting from his work on gunnery as a young naval officer in the 1914-18 war, even before going to university, Hartree's abiding interest was in partial differential equations (PDEs).

Early in his professional life Hartree became an expert on the numerical calculation of solutions to complex problems. In the early years calculation of PDEs had to be undertaken by hand, using a slide rule or a hand operated calculator. Hartree saw the great need for better means of computation and, as a professor at Manchester, in the early 1930s he started work on the construction of an analogue computer – the differential analyser. Eventually he built a full scale machine in 1935. This was used by Hartree and others for the solution of quite a wide variety of problems and could be claimed as something of a first as a computer centre for scientists. It was to become a model of what Wilkes, with Hartree's help, achieved with the EDSAC fifteen years later. One amusing example was when the first German magnetic mine was discovered. A team from the admiralty travelled to Manchester to use the machine to solve problems regarding its firing. But Hartree was too quick

for them. He suggested that they go off and get some breakfast. By the time they returned he had solved their problem by hand.

The differential analyser was a machine which undertook integration by making a wheel roll on a rotating disc, while its distance from the centre of the disc is varying. By combining such units in appropriate ways the analyser could provide solutions of complex differential equations.

As an expert in analogue computing, as well as numerical analysis, Hartree was an obvious candidate to become involved in the digital field once it came in sight. In April 1944 a committee which included Hartree recommended that a mathematical section be set up within the National Physical Laboratory (NPL). In October this recommendation was put into effect with its first two objectives being the investigation of the possible adaptation of automatic telephone equipment to scientific equipment and the development of electronic computing devices suitable for rapid computing. One suspects that some members already knew of Colossus.

Womersley (Turing's *bête noire*) was the first Director. In February 1945 he went on a two month tour of computing installations in the USA, including visiting ENIAC (still not complete). He became acquainted with the famous draft EDVAC report. About two months later Hartree also went over to see ENIAC, not then publicly known.

In the summer of 1946, Hartree spent further time on ENIAC at the invitation of the US Military in order to experience using the machine and consider its application to a wider field of scientific problems. He also himself programmed ENIAC to work on a problem concerned with laminar flow of compressible fluids. He thus became the first person outside the USA to program a problem – and a real one at that.

About the same time the NPL resolved to authorise the construction of Turing's Pilot ACE.

On 7th November 1946 the Daily Telegraph, having interviewed Hartree, quoted him as saying:

The implications of the machine are so vast that we cannot conceive how they will affect our civilisation. Here you have something which is making one field of human activity 1,000 times faster. In the field of transportation, the equivalent to ACE would be the ability to travel from London to Cambridge ... in five seconds as a regular thing. It is almost unimaginable.

But Hartree's helping hand on the birth of ACE was only the first of four early British computers with which he was concerned. Going back to February 1946, Newman (who of course had been largely involved in the

Colossus project) at Manchester submitted an application to the Royal Society for funds to start the task of building a general purpose computer. The Royal Society referred the request to Hartree and Darwin (grandson of Charles Darwin) to advise them. Hartree recommended the grant but Darwin, who was head of NPL, opposed it on the grounds that the ACE at NPL would be sufficient to serve the needs of the country. But Hartree's view won the day. As is well known, F C Williams led the Manchester project which then moved rapidly on. It is interesting to note that Williams had worked with Hartree before the war when they jointly constructed the automatic curve tracer for Hartree's differential analyser.

In August 1946 Hartree's appointment to the Cambridge Chair of Theoretical Physics was announced. In October he gave an inaugural lecture entitled "Calculating Machines: Recent and Prospective Developments and their impact on Mathematical Physics". This described ENIAC and the work that Hartree had done on it. As a 17 year old schoolboy, the author remembers purchasing the published version of the lecture in Foyles. Admittedly the contents conveyed to an A level student more excitement than understanding! To many, it was probably the first introduction to the idea of an electronic computer and what was involved in programming it. For instance Hartree describes not having realised that, if faced by having to divide by zero, the machine would not react in the same way as a human computer who would undoubtedly go back to someone and ask what to do.

Even in 1946, two years before stored programming electronic computing became a reality, Hartree saw the need for the use of sub-routines. His inaugural lecture ended with a look at what computers might do. He said:

..there are, I understand many problems of economic, medical and sociological interest and importance awaiting study which at present cannot be undertaken because of the formidable load of computing involved.

And at the time of his inaugural lecture Hartree is quoted as saying:

It may well be that the high speed digital computer will have as great an influence on civilisation as the advent of nuclear power.

But Hartree's most important contribution in the computer field was in connection with the computer that was built at his alma mater.

In 1937 Maurice Wilkes, who was Director of the Mathematical Laboratory at Cambridge, visited Manchester to see Hartree's analogue computer. Wilkes describes the occasion:

Meeting ... Professor Hartree who in later years became my close counsellor and friend. On this occasion, not only did he invite me to stay at his home, but he met me at the station and drove me to the university ... I sense that here was a level of professionalism in computing far beyond anything I had encountered up to that time.

At the end of 1945 or very early in 1946 Hartree briefed Wilkes on developments in computing in the USA. Wilkes received an invitation from the Moore School (the builders of ENIAC) to attend a course on electronic computers. Before leaving for this, Hartree was able to brief him more fully on ENIAC. It was on the boat home that Wilkes planned the original design of EDSAC, which was to become operational in May 1949. Hartree worked closely with Wilkes in developing use of the machine for a wide range of problems and, most importantly, showed users from a number of areas in the university how they could use it in their research work.

It is fascinating to read Hartree's piece "High-speed automatic digital machines and numerical analysis" which is part of his book *Calculating Instruments and Machines*, completed in May 1949. In this he covers algebraic linear equations, ordinary differential equations with 1- and 2- point boundary conditions, and partial differential equations of the different types. All these were dealt with on the basis of how the computer method of solution would differ from the old manual one. And all this before any electronic digital computer, other than ENIAC, had done any useful work.

Finally, as regards UK computers, in early 1947 the great catering firm of J Lyons & Co in London heard of the ENIAC and sent a small team in the summer of that year to study what was happening because they felt that these new computers might be of assistance in the huge amount of administrative and accounting work which the firm had to do. The team met with Goldstine at the Institute of Advanced Study in Princeton who wrote to Hartree telling him of their search. As soon as he received this, Hartree wrote and invited representatives of Lyons to come to Cambridge for a meeting with him and Wilkes. The story of LEO, the commercial version of EDSAC developed by Lyons, is well known. After Hartree's death, the headquarters of LEO Computers was renamed Hartree House. This illustrates the extent to which Lyons felt that Hartree had contributed to their new venture.

Thus Hartree was associated with the first four computers developed for scientific work as well as the first designed for commercial work. Hartree was hugely knowledgeable and hugely helpful to everyone. His unselfish and unstinting work did much to get the computer revolution

underway in the UK. Incidentally, it is to Hartree that we owe the Americanised spelling of “program”.

As indicated, Douglas Hartree was a man of huge generosity of spirit. Always willing to break off his own research to help others without any wish for acknowledgement of his help; refusing to have his name added to some of the papers which his students published; and always willing to be concerned with the most junior undergraduates. Asked how long it took him to prepare a single lecture, Hartree replied “if there are 50 students attending, it wouldn’t be unreasonable to spend up to 50 hours. As Maurice Wilkes wrote:

His knowledge of numerical analysis and, more particularly, his wide experience of computing of the most practical kind qualified him to play a major role at that critical moment. He had personal qualities of no less value. He was entirely without any sense of his own importance and could work, seemingly on equal terms, with those much younger than himself. He never attempted to lay down the law and detested those who did ... It was, in no small measure, due to Hartree that computer applications in Cambridge got off to a good start.

Though he had a very wide range of friendships in the scientific community, Hartree was not well appreciated in his lifetime by all academics. On one occasion an eminent mathematician, introducing his lecture, referred to differential equations – the subject on which Hartree was to talk – as “a sordid subject”.

Hartree’s contributions in atomic physics, in fields as diverse as the development of radar and the control of industrial processes, and in being one of the foremost numerical analysts of his time would each have earned him a serious reputation. But Hartree’s significance in the use of ENIAC, and the birth of the ACE, the early Manchester computers, EDSAC and LEO was surely unequalled in the computer field.

And perhaps his foremost achievement was that, at a time when numerical analysis was seen as a sordid subject not worthwhile for scholars, he saw that this was something that really mattered, and would be important in the future, and was able to advance the subject beyond all recognition through his early computer work.

Should Hartree’s name not be more highly honoured among those of the other computer pioneers?

KDF9 Time Sharing: Eldon 2 is not EGDON!

David Holdsworth

In the 1960s software still existed as an extra that manufacturers felt obliged to provide to enable their hardware to be used efficiently. It was quite the norm for many institutions, particularly universities, to believe that they could much improve the use of their mainframe computers by writing their own operating software, or at least making extensive modifications to the supplied software. The Electronic Computing Laboratory at Leeds University was one such institution, as was UKAEA's Culham Laboratory. Here is the story of how the Leeds KDF9 provided multi-access and batch jobs (and latterly a cafeteria) on 192 Kbytes of RAM and two 24 Mbyte discs. Yet only two KDF9s used Leeds's Eldon2 while all the rest used Culham's EGDON.

The Flowers Report in 1964 recommended that universities needed more beef in their computing, and more specifically that those universities which had installed KDF9s should have them upgraded to 32K words (192K bytes), a big (physically) fixed disc holding 4M words (24M bytes), and also a card reader, card punch and a full house of tape decks (eight, I think). This was known as the EGDON configuration, that of the UKAEA's KDF9 at Winfrith Heath (believed to be the Egdon

Heath of Thomas Hardy's Wessex) where they had developed the EGDON system that made a KDF9 look very like an IBM 7090. I recall that its cabinets were a two-tone blue design, a slighter lighter hue than IBM's blue.



As part of the deal with English Electric there were more or less national level discussions on software provision, which proved fairly

inconclusive. The system supplied by EE (English Electric) was the Time-Sharing Director, and the POST and PROMPT systems for file management and compilers for Algol 60. The EGDON system included a compiler for EGTRAN, one of the many dialects of FORTRAN that were in existence in the 1960s before the advent of FORTRAN IV and ASA FORTRAN. EGDON had just a single job stream, with input from and output to magnetic tape, with input and output “spooled” onto and off the tapes as separate concurrent operations.

All universities except Leeds moved over to the EGDON system, usually after a period of alternation of systems.

When I arrived at Leeds in October 1967, there were some experiments in trying to make the English Electric PROMPT system drive interactive teletypes. PROMPT was EE’s system for having program source and binaries held on the disc drive, an extension from the earlier POST system that held programs on magnetic tape.

After a few full and frank discussions we decided that the emphasis should be on the multi-access system, and that its load on the overall system should be so small in terms of memory occupancy and CPU usage that it could be operational throughout waking hours. EGDON seemed to be a step in the wrong direction, wasting KDF9’s elegant timesharing hardware and not well suited to the newly-emerging style of interactive computing.

So we ploughed our own furrow. The earlier experiments had been called Eldon both as a joke because of the similarity to EGDON and in recognition of the fact that the Leeds KDF9 had originally been housed in a disused chapel, called Eldon Chapel. There was also a nearby pub called Eldon. So our new venture was called Eldon 2. Around this time EGDON also underwent a revision and became EGDON 2. In retrospect, our choice of the name “Eldon 2” was a bad one, as imperfect hearing and/or imperfect diction often blurred the distinction between the two systems.

EE’s system supervisor/kernel/control program was known as the Time-Sharing Director. It allocated physical discs to programs. This was actually a recipe for maximising head movement – always non-optimal, especially so when seek time was often 300mS. Also it was a rather coarse unit of allocation. So an early part of the project involved rewriting the disc handling routines in director. EE had a rather natty algorithm for manipulating bit patterns using KDF9’s floating point instructions, and we enthusiastically embraced this technique for manipulating the bit pattern that marked allocated disc blocks. However,

many months later we spent a lot of time testing the disc before discovering a pattern-sensitive floating point hardware fault.

Most of the system was written by about five of us. Recollections are a bit vague as to who did what, but I've tried to give credit where it is due.

The teletypes were connected to a PDP-8 mini-computer (about as mini as a small wardrobe), which was also part of the Flowers upgrade. Mike Wells had designed a purely hardware multiplexer, but Peter Poole of UKAEA Culham suggested the use of the PDP-8, with 4k of 12-bit words, and a very simple design of multiplexer, leaving all the clever stuff to software – and it needed to be clever software. The multiplexer was completely unaware of characters, and just allowed the software to have control of the voltage output to each teletype and to read the voltage on each input line – nothing else. Reading a character involved sensing the individual bits at the appropriate time intervals. Transmitting a character was only slightly easier, requiring each bit to be individually sent to each teletype. Mike Wells programmed the PDP-8, beginning by writing an assembler to run on the KDF9. On start-up his program loaded the message “SYSTEM RESTARTED” into every teletype’s output buffer. It has to be admitted that the Eldon 2 system would crash occasionally. The terminal room would fall eerily silent as all the teletypes stopped. Sometimes the recovery code did what everyone hoped, and the teletypes would come back to life one-by-one. On other occasions the collective groan of disappointment was drowned out by the simultaneous typing of about 30 copies of “SYSTEM RESTARTED”.

The PDP-8’s interface to the KDF9 was much more autonomous, and proceeded as a Direct Memory Access (DMA) transfer, with two PDP-8 words controlling the transfer count and the address to which the next word was to go. This enabled some elegant and highly efficient programming. In order to load the PDP-8, the operators had to key in the bootstrap program in binary, all 3-words of it. It consisted of a zero in the counter word, all-ones in the address word, and a jump to itself in the last word of store, which was then obeyed. The KDF9 then claimed the interface and sent 1024 of its words, to comprise the 4096 PDP-8 words, the last of which overwrote the jump instruction that the machine was obeying with a jump to start things rolling.

Each line of teletype input was assembled in the PDP-8, converted to Algol Basic Symbols and then queued awaiting the next read from the KDF9. The teletype process was then queued, waiting for the last word of its buffer to become zero. The eventual reply from the KDF9 was written

over the two words that controlled the DMA transfer. The message received earlier by the KDF9 included a header that was echoed in this reply, and over-wrote the two DMA words, so as to send the rest of the reply to the buffer corresponding to the teletype for which the reply was destined. The last 12 bits were zero, and so the teletype process woke up and went through the laborious process of typing out the message to the user while the KDF9 got on with processing other people's messages. Of course, not every sequence of end-user typing constituted a valid sequence of Algol Basic Symbols. The PDP-8 software included a natty interactive facility for dealing with typos. In retrospect, we were glad that the Culham solution for driving teletypes was adopted. We were never of the same opinion regarding the multi-access system.

Inside the KDF9's 32,768 words (192 Kbytes) of store, just over 3000 were occupied by the director, and during the day another 3000 words or so were occupied by the multi-access system, which was the only program talking to the PDP-8. Within the multi-access system was an area for re-entrant segments overlaid from disc, and a 760-word buffer area into which data for each terminal was swapped in turn.

Most users spent most of their time interacting with the file editor which was written by Tony McCann. In order to save precious store, the editor made one pass through the file, and could easily be sent back to the beginning to pass through again. The disc blocks of 640 words were read into the top end of the 760-word data area, and the output blocks were constructed starting near the bottom of the 760-word area. Much of the time there was not enough extra material to fill the gap, so the output buffer grew as the input buffer shrank. It was necessary to deal with the case when the gap disappeared. Tony came to the view that we should have allowed him two separate 640-word buffers, and made the multi-access system occupy more store.

A user's program was run by putting an entry in the foreground job queue. There was one slot for each teletype, and when the entry had been written, the message "IN Q" was sent to the teletype, which then stopped until the job had run to completion. The output from the job was written into the 760 word data area on disc dedicated to that particular teletype, and was copied back to the teletype. Foreground jobs had a 30 second time limit, and normally a user could expect to wait only a very few minutes to see the output. There was also a background job queue (96 slots I think) into which multi-access users could put entries.

The initial system with on-line file editing and foreground job queue went live to end-users in the spring of 1968. The number of teletypes

soon increased to 32. It has to be admitted that 32 active users did not get adequate response, and usage rarely went above 27 or 28, as people got fed up and logged out when more were logged on.

The overall control of the machine was in the hands of the Director, and the JOBORGANISER program, with some help also from the operators. Jobs in the background queue were also selected by JOBORGANISER. I think that only one foreground job ran at once, and certainly only one job could be swapped out.

The extensive modification of the Director and the writing of the JOBORGANISER (aka JO) were the work of the author. The KDF9 had four hardware program levels, i.e. four sets of central registers, and some interlock hardware. Whenever some more store became available (i.e. a program ended or shrank its requirements), JO was loaded into all the available store. It then selected which job to run, shrank its store requirement appropriately, and turned itself into the new job. It also had the power to swap out a background job in order to make way for foreground work, and could choose to reinstate a swapped out job, by reading it back in and telling Director to reconnect it.

The KDF9 disc drive was not particularly high performance, either in terms of speed (300mS seek) or capacity (24 Mbytes). The KDF9 interface hardware (but not the Director) had the capability to drive four such discs. However, a disc stood about 5 feet high, and occupied floor space about 6 feet by 4 feet. The Leeds KDF9 did get a second disc, from a KDF9 being decommissioned elsewhere, and after a bit of local enhancement to the Director, it worked very well. Each platter had an independent head. The disc assembly had big windows, and during the daytime, with the multi-access system running all 32 heads (16 on each disc) could be clearly seen in action.

In the 1960s and early 70s, the main job of the operating system was to use the machine efficiently. Making life easy for the users (who still saw themselves as the privileged elite that they were) was a secondary consideration. Once a user's program had been allowed into the store, the priority was to get it out again ASAP. On no account let it sit there while it drives the printer, or even worse the paper tape punch. All output was written to magnetic tape – much quicker than the disc because the output blocks were just written in the order that programs generated them – and sorted out afterwards.

This sorting out was done by a program called the Destreamer, because it took the streams of output and untangled them onto the printer or punch. Colin Harman wrote the code of the Destreamer. Colin insisted

on writing Usercode (the official EE assembly language) without spaces between the instructions (which are allowed several to a line). Others of us never persuaded him to change his style, even though we all found it difficult to read. He certainly knew that his style would not reduce the size of the binary program. KDF9 allowed user programs to own and drive peripherals directly. Furthermore, there was a curious instruction (INTQq) which let other programs run until one of your own transfers ended, and then switched back to you in a very efficient manner. We could not afford to use one of the four levels for this printing activity which ran almost non-stop, so it was embedded in the multi-access system, exploiting this quirky instruction to great effect.

The result of the combination of limited disc space and effective multi-access was that the file space soon filled up. Tony McCann wrote software that combined the role of back-up with the management of off-line files. We had a two-level filestore delivering end-user service long before George3, and with tape-to-tape processing from the outset.

The KDF9 systems had originally firmly nailed their colours to the Algol 60 flagpole, but there was a steadily increasing pressure to run FORTRAN. So we ran bits of EGDON overnight, but really we needed a FORTRAN compiler, so it was



either EGDON or write one. Unlike other universities we chose the latter course. David Johnson and I worked together on this, writing our compiler in our new assembly language KAL4, which had been implemented as an MSc project by John Thomason. On KDF9, floating-point arithmetic with integers is exact up to 2^{39} . So to avoid the hassle of type-checking, all variables were stored in floating-point. KDF9 also had an arithmetic stack of 16 cells (the nesting store) instead of more traditional general registers. We decided not to bother worrying about this limit, and generate code that would fail at run time if someone wrote a very heavily bracketed expression. (I believe that it happened just once.)

We generated KAL4, and pushed this at our new assembler. KAL4 had mnemonic identifiers, so the names from the FORTRAN program were preserved, and the tables output from the assembler could make some diagnostic sense. The parser was hand coded, and each statement was compiled in its entirety. It was workman-like generated code, without optimisation. The KDF9 register architecture was a very good match to arithmetic computation.

Without this FORTRAN compiler, we would probably have had to abandon Eldon 2 in favour of EGDON, as it became clear that FORTRAN was doing a VHS to Algol's Betamax. Instead, we were able to abandon EGDON, which by this time had an Algol compiler.

So what of Algol?

We certainly made a load-and-go Algol system that was very heavily used by teletype users. I cannot remember who implemented it. I think that the late Richard Hughes had an important role and probably also David Johnson. This used EE's Whetstone Algol, known colloquially as WALGOL. The implementation translated Algol into pseudo-code and then executed that by interpretation. Quite like the initial Java implementations.

The only way to get good execution speed of an Algol program was to compile it with the Kidsgrove compiler. This multi-pass compiler took rather a long time, too long for the foreground job queue. For a long time, KALGOL (as it was colloquially known) was seen as magical, and beyond tampering. Several passes (at least one of them backwards) compiled the Algol into Usercode, which was then combined with the run-time library code, all in Usercode. The whole lot then went through EE's Usercode Compiler, a misnamed and rather slow assembler. We discovered that this Usercode stage was more than half the compilation time, for smallish programs, much more. So we bundled all the commonly used library routines into a lump with a vector of jumps to all the entry points at the start, and assembled it. We made a new library in which these commonly used routines were replaced by jumps into the vector, then arranged to start the generated code part way up the address space, and bingo: KALGOL compilations were quick enough to run as foreground jobs. Also we had a version of the library for running a program from the job queue, and one for operating traditionally with paper tape data. Both operated with the same binary program.

And then there was BASIC.

NPL implemented a BASIC compiler (the work of Tony Hillman, I believe), so we had a RUN BASIC command alongside RUN WALGOL and RUN FORTRAN.

We also had graphics facilities thanks to Rae Earnshaw, and even one Tektronix storage tube alongside one of the teletypes.

In 1972, George3 arrived in the shape of a 1906A. It was now becoming accepted that computers were not just for post-graduates and staff, and that under-graduates might also get a look in. So the KDF9 became a student facility. The machine was reconfigured, losing a couple tape decks to reduce floor occupancy, and gaining a second printer in the terminal area. A rather cramped ante-room housed the original printer, a paper-tape reader and a card reader. The machine now operated as an open-shop cafeteria for students. I implemented a facility in Director to hold a read on both the paper tape reader and the card reader. We disabled the hooter that KDF9 used to tell operators that a peripheral needed their attention (e.g. to load a paper tape). When the paper tape read was answered, WALGOL was loaded with the device number of the reader in the nesting store. There was an OUT (system call instruction) to get the contents of the buffer. A read on the card reader loaded our FORTRAN compiler. All the while the multi-access system was also running upstairs.

So at the end our configuration was supporting approx 25 teletype users, and a cafeteria offering FORTRAN and Algol, all with two disc drives (24 Mbytes each), 192 Kbytes of core store, five or six magnetic tape decks.

So what did we do wrong? Why did every other university switch to EGDON?

We did not have enough confidence to sell the system before we had implemented it. We chose a name so like EGDON that we frequently had to explain that it was not a mispronunciation or typo. We over-estimated the difficulty of compiling FORTRAN for the KDF9 architecture, and only started FORTRAN implementation around 1970. EGDON had Algol before we had FORTRAN. If we weren't sure that we could make it all work, it was not surprising that others had already taken their decisions before Eldon 2 could offer a full menu.

So what did we do right? Why did Eldon 2 outlive EGDON?

The only other installation to use Eldon 2 was NPL, whose two KDF9s were shutdown on 29 August 1980. The Leeds machine ran until about 1976. By exploiting the architectural features of KDF9, the mature Eldon 2 gave better throughput than EGDON, especially under heavy

multi-access load, which by 1970 was becoming the expectation of users. Instead of playing catch-up by emulating somebody else's operating system, we were focussing on the then new-fangled multi-access, but still with the emphasis on efficient machine utilisation, which was of serious concern when computers were expensive to buy and to run.

We did actually write it up, and it appeared in Computer Journal - sw.ccs.bcs.org/KDF9/CompJ69.html.

So what did we do next?

The ICL1900 came next with George3. Although it offered facilities for interaction not available with Eldon 2, it did not support as many teletypes when first installed. We implemented Eldon 3 (written in Algol 68-R) to drive terminals with a hope of better facilities and more throughput. It worked, but not as well as we hoped. Then came the Amdahl with VM/CMS, before the world saw the light with UNIX and the networked world of client/server – oh yes, and PCs.

I would like to acknowledge helpful comments from Mike Wells.

Editor's note: David Holdsworth first met KDF9 as a user in Oxford in 1965 modelling quark calculations, but also showing an interest in the system's software, where he made one or two improvements (changes?). He worked at Leeds University in Computing/IT from 1967 to retirement in 2004. He now lives in the Yorkshire Dales, where he still involves himself in software conservation. He believes that software is only properly conserved if you can run it. David can be contacted at ecldh@leeds.ac.uk.

CCS Website Information

The Society has its own website, which is located at www.computerconservationsociety.org. It contains news items, details of forthcoming events, and also electronic copies of all past issues of *Resurrection*, in both HTML and PDF formats, which can be downloaded for printing. We also have an FTP site at [ftp.cs.man.ac.uk/pub/CCS-Archive](ftp://ftp.cs.man.ac.uk/pub/CCS-Archive), where there is other material for downloading including simulators for historic machines. Please note that the latter URL is case-sensitive.

Elliott Computers for Motorway Signalling

Richard Burwood

Half a century ago, as the new motorway network was being constructed across Britain, no provision was made for communication with motorists. This proved to be a tragic mistake as numerous fog-related accidents were to demonstrate. A system of hazard warning signs based on winking car headlights was hastily improvised, but each one had to be operated manually, limiting their responsiveness to events, and each able only to indicate an unspecified condition. Obviously something better was required. Thoughts turned to a system of centrally-controlled dot-matrix signs linked into a computer allowing a whole stretch of motorway to be controlled from a single place with immediate response and with a range of indications available. This is essentially the system still in use today albeit the original hardware has long since been replaced.



The Ministry of Transport ordered a number of computers from Elliott to be used to control traffic signals on motorways in England and Wales. A total of seven (12 processors) were supplied from 1968 to 1972. The sites were:-

- CP1 North West – near Preston
- CP2 North East – between Sheffield and Wakefield
- CP3 Midlands – northern suburbs of Birmingham
- CP4 South West and Wales – near Bristol
- CP5 East – between Mill Hill and Elstree
- CP6 South – M3 Hampshire
- CP7 London West – near Heathrow

Prototype

The prototype system was CP7, to control the section of M4 from Hammersmith to Heathrow. This was an “ARCH 9000” (Articulated

Control Hierarchy) containing a single Elliott 903 processor. There was one memory module only (8K x 18 bits). This computer seems to have been delivered in December 1968. *Computer Survey* vol.9 no.6 described this as a Marconi-Elliott 903 for the M4 Metropolitan Motorway.

The teletype used to operate the system was modified by inscribing some key caps with symbols such as “STOP” and “30”. The control software printed the full text.

The software team included contract staff from Vaughan Programming Services.

The executable software was copied onto Mylar paper tape at some point. This lasted a considerable time, but eventually Vaughan had to start making copies of the system tape on normal paper tape.

The non-standard prototype outstation equipment remained in service until replaced by a new system.

Pre-production

The second system was CP2, to control the M1, M18, and A1(M) around Sheffield and Wakefield. This was an ARCH 9000 with a single 903 processor, and again seems to have had only 8K memory. This probably had a standard teletype but only for engineering purposes - the Control Office probably had a Texas Silent 700.

Outstation equipment was different, perhaps most particularly in that levels were inverted. The modem frequencies for “1” and “0” were exchanged.

Computer Survey vol.9 no.6 described this as a Marconi-Elliott 903 for the West Riding Motorway.

Vaughan modified the software to handle the production responder type and the pre-production type. The Wakefield section of motorway was transferred to CP1 later, but Sheffield remained on CP2 until that was closed down.

Production Systems

Four further systems were ordered for 1971 and a fifth for 1972. These were described in *Computer Survey* vol.9 no.6 as GEC-Elliott MARCH 9050, the latest marketing name at that time.

The first four were delivered in 1971 or 1972 (records are unclear), for CP1, CP4, CP5, and CP6. The fifth, for CP3, was a year later (1972 or 1973).

“GEC” labels were stuck over the “EA” insignia on the consoles. These labels had a remarkable tendency to disappear over the years.

These had a dual processor and dual ARCH 9000 - not 9050. Perhaps the 9050 was more expensive or in short supply, or 9000s were going spare. A standard 9000 would have had level convertors for the 903 processor, but these had different ones for 905 processors. The standard 905 cabinet had no memory - the second processor occupied that position. The cabinet was sandwiched between memory cabinets for the two half-systems. Each memory cabinet had two modules each by Ampex and of 16K x 20 bits, of which 19 were used for the 18-bit words and for parity. Thus each half-system had 32K memory initially.

CP6 never became operational, its potential workload being transferred to CP5 and CP4. CP6 in Hampshire was used by Vaughan staff (from Hertfordshire) for software development and testing purposes. Parts of the equipment were removed to augment other sites, as 32K was found too limiting, and more interfaces and modems were needed for more Control Offices. At some point, the remains of CP6 were transferred to Vaughan’s office in Ware, Hertfordshire, so that I could do development work in house.

Provision was made for the potential boxes of motorways around Manchester and Birmingham. The double 32K from CP3 were transferred to CP1, increasing that to 64K per processor. New memory was fitted at CP3, where each half-system was given a single memory cabinet with two 32K modules. CP4 remained 32K, and part of the memory from CP6 was transferred to CP5 to increase that to 48K.

Later, it was decided that a 40K or 48K version of the software should be used at all sites. CP5 had 48K, CP3 had 64K, and CP1 had 64K, but CP4 had 32K only. A new computer system near the junction of the M6 and the M42 was gradually taking over Control Offices from the 905 systems, and it was realised that the timescale for that and the timescale for the new motorways around Manchester meant that CP1 did not need its 64K, so that was reduced to 48K and the spare memories taken to CP4 to raise its size to 48K. CP3 retained 64K.

When CP1 was shut down, the equipment was taken to Ware and the computer room re-assigned to another new system, as the new system by the M42 would be above maximum loading if trying to run the whole country (England and Wales). When this became operational, it took over the North West, the North East, and the East Midlands. The equipment from CP1 plus that already at Ware (from CP6) was reconfigured so I would have a fully dual system with 64K on each side, and with four

processors - each console cabinet had one operational processor rack and a spare rack. I used the extra type-3 power supplies to ease the loading on the mains distribution. When either half-system was switched on, the load from 4 x 16K memories, paper tape station (the punch switched on initially and then switched off), and the half-ARCH produced a surge current exceeding 13A.

Production Software

The intention was to use standard software at every site. Only the database would vary between sites. This seems to have worked well at first. Software was known by the generic code of 4010, with the GEC software being type 4011, 32K, issues 1 to 6.

How did the software cope with different memory sizes? It didn't at first. Vaughan modifications of the 32K software (without "Lane Divert"), designated 4012, continued for CP1 and CP4 (issues 7 and 12). It is not clear which had issues 8 and 10 - or were these for CP3 and CP5? Both CP1 and CP4 received a merged issue 14 - not a correct issue number, as the 4013 had used 14 by then. Issue numbers applied to the 4010, not to specific versions of it (except perhaps for that issue 14).

Vaughan modifications (with "Lane Divert") were designated 4013. Issues 11 and 14 were for CP3 and CP5. It is not clear which (one or both) had issues 15 to 22, nor whether the revision to use 40K for the software started at issue 20, 21, or 22. 19 was 32K, 22 was 40K. Issue 23 and 23.1 were 40K, for both CP3 and CP5. Issues 23.2 and 23.3, also 40K, were for all sites. I worked on them and produced also issue 24 (40K) for all sites, but this was not installed.

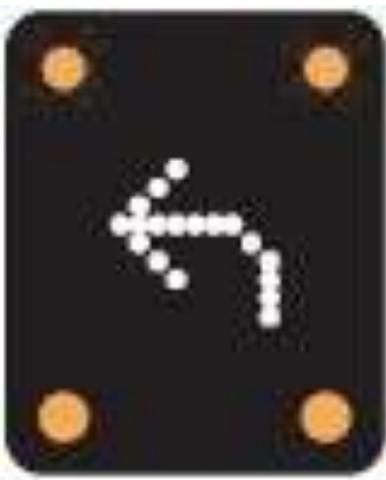
Vaughan persuaded the Department of Transport to allow us to rationalise the software, introducing our MaCE design, invented by Dina St. Johnston, which is almost certainly the best concept for any continuous control system. Unlike a Real Time Operating System, which overloads itself with housekeeping work as it tries to keep track of activities, MaCE is most efficient at maximum workload.

Version 4014 was designed for 48K, and issues 25 and 26 were introduced before the West Midlands decided to fund modification of the software for Tidal Flow. Version 4015 used the first 32K for code and running data, and the next 16K for Site Data - issues 27 and 27.1. Issue 28 allowed the Site Data to expand over 16K if the target system had 64K.

CP1, CP4, and CP5 were closed down over the period from 1980, and most Control Offices were transferred from CP3, but my unique

Tidal Flow software kept CP3 controlling the Birmingham Control Office until August 1994.

Site Data



Site Data was installed on separate paper tapes, and was specific to the site. This was prepared by Vaughan, taken to CP6 for conversion to binary, and copies produced for issue. When CP6 was shut down, the standby half-system of CP6 was used instead - I used that a number of times from 1978 until the development system became available in the Vaughan office.

Most site data work was performed by another member of staff at first, but as I was better able to vet the data changes for errors, I took over that work. The data preparation tool (known as SETUP) was re-written for issue 25 and enhanced for issue 28 to use 64K for CP3 or 48K for other sites.

Bletchley Park

The 905 taken to the National Museum of Computing in 2008 was part of the Vaughan system. This had been stored by Andrew St.Johnston at his farm since 2004. It is unfortunate that important spares for the 905 - and the dual ARCH - went for scrap. I still have the VDT that replaced the Teletype, and many of the paper tapes and documentation, including listings and some source tapes.

Vaughan's other half system was taken by someone else who wanted a working 905.



Contact details

Readers wishing to contact the Editor may do so by email to dik@leatherdale.net, or by post to 124 Stanley Road, Teddington, TW11 8TX. Queries about all other CCS matters should be addressed to the Secretary, Kevin Murrell, at kevin@ps8.co.uk, or by post to 25 Comet Close, Ash Vale, Aldershot, Hants GU12 5SG.

That Atlas Switch Panel!

Dik Leatherdale

Readers will remember that in *Resurrection 48* we asked whether anybody could identify a switch panel, known to be from the London University Atlas 1 and labelled “CORE STORE”. Well, as you might expect, answers came there several.

Mike Robinson was first out of the traps with the ink barely dry. He recalls working on the development of the ICL 1906A with a similar panel controlling Fabritek core store. The 1906A had double word memory management and so, at least at core store level, was effectively the same 48-bit word length as Atlas.

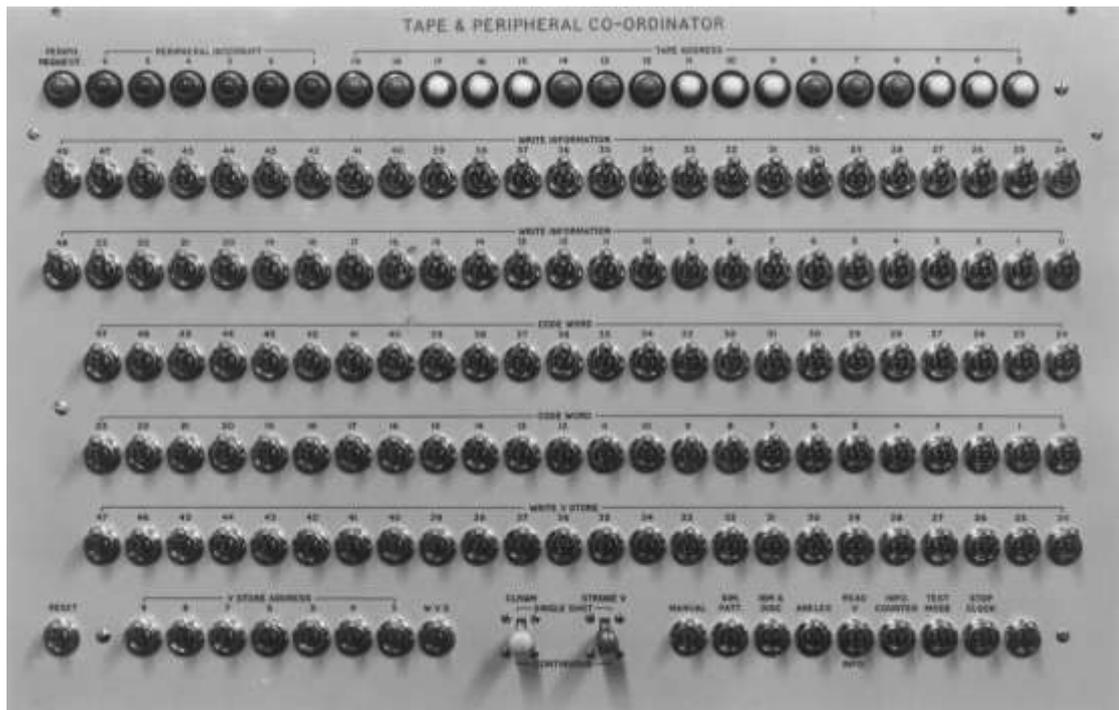
Then Chris Burton chimed in with several photographs taken during Atlas commissioning at West Gorton. They show panels similar in general appearance but obviously built for varying purposes. The other thing that they seem to have in common is that each was each was mounted on a trolley underneath an oscilloscope.

I consulted Brian Hughes – the Field Engineering Supervisor on the London Atlas. Brian didn’t recognise the mystery panel, but remembered the trolleys which were used to wheel the oscilloscopes around. The switch panels and associated test gear was not much used in the field it seems.

What are we to make of all this? It would appear that the panels were not part of the machine itself but its associated test equipment. All were made up to a common style (all by the same person at West Gorton perhaps) and from the same components, but in different configurations serving different purposes. It’s also possible that the equipment was reused during the development of the ICL 1906A a few years later.

Not quite there then yet, but we know a lot more than we did when we started. Thanks are due to everybody who contributed. If anybody else has had their memory jogged, please let us know.

The photographs are below.



A different panel but in the same style and using the same components



This photo is taken from the www.chilton-computing.org.uk website and is labelled "Engineers at work". It shows yet another design of panel mounted on a 'scope trolley. Courtesy Rutherford Appleton Laboratory.



Atlas being commissioned showing no less than three commissioning trolleys!



Close up of one of the commissioning trolleys with a fourth type of test equipment switch panel.

System 25 – ICL’s Cinderella Success of the Eighties

Michael Knight

We all know the basics of the traditional Cinderella story. In this more recent version, with System 25 in the eponymous role, some veteran ICL readers may want to choose their own actors for the other characters.



ICL’s System 25 was developed as a program-compatible successor to System Ten, which was acquired from Singer Business Machines in 1976 (see *Resurrection 47*). The development was code-named System 29, which, when leaked, prompted journalistic speculation that it was some kind of “swing-machine”. In reality, ICL just had a fixation about “29”. How did 29 become, publicly, 25? Well, it was sort of between DRS20 and ME29....

Compatibility with System Ten would sustain that product’s large, profitable, international customer base and the very large third-party application software investment in it. System 25’s architecture and technology, however, were radically new, based on a variety of microprocessors networked by three parallel highways, an approach which anticipated, for example, DEC’s VAX11 evolution from a single CPU and highway by several years. The “instruction processor”, emulating System Ten, was a 96-bit bit-slice micro, the control processor

an 80-bit bit-slice, and Intel 8 and 16-bit micros were also deployed. 50K or 70K System Ten instructions per second were executed, dependent on model; quite enough. On some business/administrative applications, this machine could match a 1MIP VAX11-780. The emulation was remarkably successful, replicating even most of the System Ten “quirks” which application programmers were not supposed to use, but did. The architecture could, and eventually did, support other personalities, but the requirement was basically for a System Ten replacement. System 25 also pioneered in ICL use of Winchester-technology (35MB) fixed discs, sourced from Micropolis, backed up by 10MB cartridge tape drives (or 60MB reel-to-reel), which was a culture shock for the System Ten community, used to the tangible reassurance of exchangeable disc packs. In short, Cinderella was rather attractively endowed.

Launch was planned for July 1980. But delayed by development budget constraints and corporate angst, Cinderella nearly didn't get to the ball at all as ICL's financial wheels fell off. Was there money for a launch, following the lavish, theatrical launch of ME29? What about the large System Ten-120 inventory, inflated by a rash earlier senior management diktat? Could it *not* be launched, since market and media were expecting it? Discreet take-over talks occurred, and an unsympathetic Thatcher government was involved. Trades unions made the talks national news, revealing Univac as the most likely suitor, and wrongly predicting the demise of ICL's mainframes and the loss of 6,000 jobs (about half the eventual number, as it turned out). Oddly, they stressed that details of “ICL's secret new minicomputer” had been disclosed. They meant System 25. Indeed, they had been disclosed. “Due diligence” requires that an acquisitive business understands what it might buy.

Shortly after the arrival of the new Wilmot/Laidlaw top management, supported by government-guaranteed loan facilities (which guarantees were never required), System 25 was quietly launched, world-wide, in June 1981. The budget was more than usually modest. Several pre-production machines were dispersed for demonstrations. Product introduction and launch preparations were arduous, given the disappearance of secretarial support in the financial crisis and a bizarre obsession with secrecy. Voluminous keyboarding was mostly done at night, off-premises, on an elderly portable typewriter, by a patient working mother; copying, mailing, telexing, faxing, etc. were done by the small handful of relevant marketeers. The smallest model was temporarily withheld, until the System Ten inventory was cleared, which, aided by the launch and a price cut, occurred with startling rapidity.

Cinderella was at the ball – but only until midnight, which struck in October 1981, with the launch of DRS20 and the Networked Product Line (NPL).

Wilmot declared ICL had two strategic product lines: 2900/VME and DRS20; all else was “tactical” – including “doomed” System 25, the press reported gleefully. After 1985 Cinderella would be history. She was obsolescent, Wilmot remarked as Sainsbury’s were poised to sign for around 250 S25 in-store point-of-sale terminal controllers. (They did sign later, though.) She was obsolescent because she wasn’t CMOS, the chip technologist said, and her applications software was old. Actually, there were elements of CMOS in S25, and he himself had cancelled a new international accounting applications suite, in BASIC, earlier commissioned from Holland Automation, writing off £250,000, including the compiler. An allegedly ready-made alternative, from the USA, proved a predictable hoax, and so development of a new suite, Superstars, was belatedly begun, in Sydney. Eventually, Cinderella would get other ICL software in her trousseau, including COBOL, IAS (Interactive Applications Support), word-processing and a version of the venerable PROSPER spreadsheet software.

The aspiration for two strategic product lines was eminently sensible, with benefits as numerous as they were obvious. At the sub-mainframe level, it had been explored in detail in the late 1970s, under the title of the High –Volume Products (HVP) Project. Ambitious “marketing requirements” were addressed by development proposals from both Kidsgrove and Utica. This all came to a dead stop in 1978, when Ed Mack, with cunning exaggeration, advised a previously enthusiastic Peter Ellis, Worldwide Marketing Group Director, that it would mean starting “another VME/K”. Ellis was already traumatised by the cost and duration of VME/K’s development, which was Mack’s alternative 2900 Operating System. The opportunity was lost, forever. Wilmot probably never heard of HVP but one of his first decisions was to cancel VME/K. DRS20, although arguably a reflection of the HVP work, was a pale reflection, its architecture and technology not quite equal to the range of needs and roles previously envisaged.

A surreptitious Cinderella rehabilitation campaign began, using confidential revelations to the leadership of the System Ten/25 Users’ Group as a reliably porous medium. The press learned, for example, that production of System 25 was anticipated until “the end of the decade” (if only in Poona, India, where ICL’s partner ICIM primitively produced S25s under the inscrutable name of “System 101”). In fact, Letchworth production outlasted Poona’s, and the last production machine was

presented to the Science Museum in 1990. Significant sales and forthcoming developments were similarly communicated, from the kitchen closet.

Wilmot was also concerned about cost-of-sale. Despite an excellent production cost to sale price margin, he believed that each System 25 sale cost ICL money. To address this, in 1982 numerous Computerpoints, eventually over 30, were established world-wide, dedicated, supposedly, to cost-effective retailing of S25s with line-of-business “solutions” to small businesses. These were re-incarnations of the customer centres established some ten years earlier to sell 1901As and 2903s. There had been no way to apportion the huge costs of ICL’s world-wide marketing and sales group (WMG) to specific product lines. Now there was, but only for System 25, and only partially; Sainsbury-style sales would not go through Computerpoint. Cinderella was strip-searched. The performance and profitability of Computerpoints were scrutinised closely; results ranged from outstanding to dismal and, overall, they did not dispel Wilmot’s suspicions. A quiet marketing decision was made, meanwhile, to re-position and re-launch Cinderella: not just a small business system, but also a small-ish system for big business. The emphasis switched to multiple-system sales, particularly with the 9500 point-of-sale terminal range.

NPL embraced the emerging “industry standards” for Open Systems Interconnect (OSI). But Robb (Wilmot) also said that ICL should “surround IBM”, which would require conformity to IBM’s Systems Network Architecture (SNA).



(To some in ICL, “Robb said” had acquired the sanctity elsewhere of pronouncements by Jesus or Mohammed – probably far from his intention.) The outstanding S25 developers, hard and soft, soon equipped the product for both roles. As a tactical product, it was not supposed to support OSI on Ethernet local area networks, although it did, so from 1984 it talked OSI to itself, under the pseudonym of High-Speed Link (HSL). The S25 architecture had retained the System Ten limitation of 20

program partitions, which for a few prominent users with substantial System Ten “legacy” software had become a constraint. HSL, albeit very inefficiently at first, overcame this by joining two or more physical systems transparently as a logical entity with a shared filestore. As for IBM’s SNA, Cinderella adopted the relevant protocols sooner and more completely than its IBM rival, System 36 – and even helped debug them.

The re-launch events were mainly internal, world-wide, and occurred late in 1982. They were unusual; not a morning’s razzmatazz presentations followed by a moistened lunch, but three days of teaching, from architecture, pricing and competition to instruction on OSI and SNA, with sales presentation aids to match.

Meanwhile, back in Baron Hardup’s tactical kitchen, Cinderella was getting leaner and fitter. The developers took emerging opportunities to reduce the logic and printed circuit board (PCB) count, and hence the production cost. A company-wide livery and logo change gave the chance to replace the original cabinet, an imposed mainframe module engineered to resist direct tank assault. A slimmer, lighter and cheaper office-style enclosure emerged. After much argument and simulation, the imposed military-spec. standard for PCB connectors, gold alloy, was relaxed, yielding an estimated £500,000 saving on a typical year’s production of S25. A 1 megabyte memory board was developed, pioneering 256k chips in ICL, not for the memory-lean S25 applications, but as a disc file cache, producing some dramatic performance improvements. It also supported another, surreptitious development: a co-processor board for memory-hungry Unix, applications for which could use a real, proven file-store (S25’s DMF III), unlike most Unix small system offerings at the time. System Ten/25 and Unix applications could co-exist, using common files. This, apparently, was not formally released until early 1988.

A slimmed-down, rack-mounted version was prototyped, dualled with HSL for reliability. This was to provide itemised billing and other now-familiar added-value services in Plessey digital telephone exchanges. Alas, Plessey did not get the BT business. However, another slim version, for one or two users, was produced in desktop form. It looked like a PC, as Wilmot coldly remarked. What can it do? What would it sell for? “A lot more than a PC” was the answer in both cases. This, the S25/1, was announced after much argument in 1984, burdened initially with a hugely artificial transfer value uplift (more on transfer values later). Its purpose was to complement earlier, larger models, to complete an affordable range for multiple-system, multi-site customers. The worry was that it would compete with DRS20. After some months,

the transfer value was reduced, quietly but not without further, irrational acrimony.

The tactical kitchen was getting crowded. A plethora of products was arriving, mostly with little or no development by ICL and mostly not manufactured in-house, with only market diversity and hardware, software, service and support incompatibility in common; the exact opposite of the High-Volume Products vision. They included, as examples: DRS8801 (word processor), a PC (from Rair), PERQ (an American monochrome graphics workstation), CLAN (an American-sourced multi-user UNIX box), and OPD (One per Desk); a marriage of telephony and limited, idiosyncratic personal computing, embracing elements sourced from Sinclair, Psion and others. CLAN was intended to catch the perceived UNIX wave with virtually no acquisition cost, and supposedly no cost-of-sale issues, since it would be sold exclusively by dealerships. Unfortunately, the machine had actually been developed for PICK, a lean and largely forgotten operating system for small interactive systems. Use of UNIX required further hardware development by the supplier (Datamedia), notably for memory management; an unforeseen and embarrassing cost which was concealed by raiding Cinderella's modest development purse. After all, CLAN was supposed to replace her, the story went; she would "wither on the vine". In fact, the CLAN business plan was fatally, obviously flawed; it could never show a profit. But worse was to come.

Her "strategic" sisters, DRS20 and DRS 100 (introduced in 1984), were deemed not to be attracting sufficient suitors. Specifically, that product line's transfer value (TV) was too high to reconcile competitive pricing with adequate profit for sales units, so sales volumes were disappointing. This problem requires some awareness of ICL's algebra of product finance. Various uplifts, or "taxes" were added to the unit production cost of each product, chiefly "development recovery" (DR), resulting in a TV, which was effectively the product's cost to the sales unit. In aggregate, the difference between a sales unit's gross revenues, including product sales, and its own operating costs plus TVs of products sold, determined its profitability, and the prosperity or survival of its management. Low-margin products were therefore not attractive. Also in aggregate, DR was intended to fund the corporate development budget (at around 7% of turnover), but there was no particular relationship between the actual development costs of a product and the level of DR "tax" it carried.

Somehow, DRS20's development recovery tax, and thus its TV, had to be reduced, without reducing overall DR, and without further internal

taxation of mainframe products. In winter 1984-5, a study of all the sub-mainframe product lines was undertaken, to see which of them might contribute to the reduction of the DRS transfer value. There was little choice. OPD, for example, with an end-user price, via distributors such as BT, of £1195 or more, yielded ICL a gross profit of barely £30 per unit. Only one line had a sufficiently lean unit production or buy-in cost and large enough sales volume to help: the already heavily-taxed Cinderella. By then, the System Ten/25 product line had become ICL's third largest revenue earner (after mainframes and DRS), and second largest margin earner (the difference between revenues and production costs).



That was the easy part. To determine the size of Cinderella's increased burden, and DRS' reduced burden, while squaring the development recovery book, required difficult simultaneous guesses about the positive effect on DRS sales volumes and the negative effect on Cinderella's.

Curiously, when these changes had been put in place, System 25 sales experienced a bit of a surge, I was told (having just left ICL). In any event, Cinderella lived on, perhaps not happily ever after, but to a decent and dignified old age, until 1990.

Michael Knight was System Ten/25 product line marketing manager 1979-84. He can be contacted at michaelknight242@tiscali.co.uk.

Forthcoming Events

London Seminar Programme

14 Jan 2010	50 years of Advanced Programming – an Anniversary Seminar on Algol 60	Panel of Speakers
18 Feb 2010	The CDC 6600	Dik Leatherdale & John Fernbank
18 Mar 2010	WITCH & CADET	Kevin Murrell
15 Apr 2010	To be decided	
20 May 2010	Pegasus @ 50	Panel of Speakers

London meetings take place in the Director's Suite of the Science Museum, starting at 14:30. The Director's Suite entrance is in Exhibition Road, next to the exit from the tunnel from South Kensington Station, on the left as you come up the steps. Queries about London meetings should be addressed to Roger Johnson at r.johnson@bcs.org.uk, or by post to Roger at Birkbeck College, Malet Street, London WC1E 7HX.

Manchester Seminar Programme

16 Feb 2010	An Overview of Some CCS Projects	Chris Burton
16 Mar 2010	The LEO Story	Gordon Foulger

North West Group meetings take place in the Conference Room at MOSI – the Museum of Science and Industry in Manchester – usually starting at 17:30; tea is served from 17:00. Queries about Manchester meetings should go to William Gunn at william.gunn@ntlworld.com.

Details are subject to change. Members wishing to attend any meeting are advised to check the events page on the Society website at www.computerconservationsociety.org for final details which will be published in advance of each event. Details will also be published on the BCS website (in the BCS events calendar) and in the Events Diary columns of *Computing* and *Computer Weekly*.

Museums

MOSI : Demonstrations of the replica Small-Scale Experimental Machine at the Museum of Science and Industry in Manchester have been suspended due to development work in the museum. Resumption is likely to be late next summer.

Bletchley Park : daily. Guided tours and exhibitions, price £10.00, or £8.00 for concessions (children under 12, free). Exhibition of wartime code-breaking equipment and procedures, including the replica Bombe and replica Colossus, plus tours of the wartime buildings. Go to www.bletchleypark.org.uk to check details of times and special events.

The National Museum of Computing : Thursday and Saturdays from 13:00. Entry to the Museum is included in the admission price for Bletchley Park. The Museum covers the development of computing from the wartime Colossus computer to the present day and from ICL mainframes to hand-held computers. See www.tnmoc.org for more details.

Science Museum :. Pegasus “in steam” days have been suspended for the time being. Please refer to the society website for updates.

North West Group contact details

Chairman Tom Hinchliffe: Tel: 01663 765040.
Email: tom.h@dsl.pipex.com
Secretary William Gunn Tel: 01663 764997.
Email: william.gunn@ntlworld.com

Committee of the Society

Chairman: **Dr David Hartley FBCS CEng:** david.hartley@clare.cam.ac.uk
Secretary & Leader DEC Project: **Kevin Murrell:** kevin@ps8.co.uk
Treasurer: **Dan Hayton:** Daniel@newcomen.demon.co.uk
Chairman, North West Group: **Tom Hinchliffe:** tom.h@dsl.pipex.com
Editor, Resurrection: **Dik Leatherdale MBCS:** dik@leatherdale.net
Website Editor: **Alan Thomson:** alan.thomson@bcs.org
Archivist: **Hamish Carmichael FBCS:** hamishc@globalnet.co.uk
Meetings Secretary: **Dr Roger Johnson FBCS:** r.johnson@bcs.org.uk
Digital Archivist & Leader, Our Computer Heritage Project:
Professor Simon Lavington FBCS FIEE CEng: lavis@essex.ac.uk
Science Museum representative: **Dr Tilly Blyth:** tilly.blyth@nmsi.ac.uk
MOSI representative: **Catherine Rushmore:** c.rushmore@mosi.org.uk
TNA representative: **David Glover:** david.glover@nationalarchives.gov.uk
Codes and Ciphers Heritage Trust representative: **Pete Chilvers:**
pete@pchilvers.plus.com
Leader Colossus Project: **Tony Sale Hon FBCS:** tsale@qufaro.demon.co.uk
Leader, Elliott 401 Project: **Chris Burton CEng FIEE FBCS:**
cpb@envex.demon.co.uk
Leader, Bombe Rebuild Project: **John Harper Hon FBCS CEng MIEE:**
bombe@jharper.demon.co.uk
Leader, Elliott 803 Project: **Peter Onion:** Peter.Onion@btinternet.com
Leader, Pegasus Project: **Len Hewitt MBCS:** leonard.hewitt@ntlworld.com
Leader, Software Conservation Project: **Dr Dave Holdsworth CEng Hon FBCS:**
ecldh@leeds.ac.uk
Leader, 1301 Project: **Rod Brown:** sayhi-torod@shedlanz.co.uk
Leader, Harwell Dekatron Computer Project: **Tony Frazer:** tony.frazer@tnmoc.org
Dr David Anderson: cdpa@btinternet.com
Professor Martin Campbell-Kelly: M.Campbell-Kelly@warwick.ac.uk
Peter Holland: p.holland@talktalk.net
Dr Doron Swade CEng FBCS MBE: doron.swade@blueyonder.co.uk

Point of Contact

Readers who have general queries to put to the Society should address them to the Secretary (see page 28 for postal address). Members who move house should notify Kevin Murrell of their new address to ensure that they continue to receive copies of *Resurrection*. Those who are also members of the BCS should note that the CCS membership is different from the BCS list and is therefore maintained separately.

Resurrection is the bulletin of the Computer Conservation Society. Copies of the current issue are available from the Secretary for £5.00 each.

Editor – Dik Leatherdale

Typesetting – Dik Leatherdale

Cover design -- Tony Sale

Printed by the British Computer Society

© Computer Conservation Society