


Issue Number 43

Summer 2008

Computer

RESURRECTION

The Bulletin of the Computer Conservation Society



science
museum

 **BCS**
THE BRITISH COMPUTER SOCIETY

THE MUSEUM
OF SCIENCE &
INDUSTRY
IN MANCHESTER

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 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,
- ◇ Develop awareness of the importance of historic computers,
- ◇ Encourage research on historic computers and their impact on society.

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 Working Parties 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 43

Summer 2008

Contents

Greetings!	2
<i>Dik Leatherdale</i>	
Chairman's Report for 2007-08	2
<i>David Hartley</i>	
News Round-Up	4
Society Activity	9
Pioneer Profiles – Christopher Strachey	11
<i>David Barron</i>	
Computers at Metrovick : The MV950 and AEI 1010	14
<i>Ron Foulkes</i>	
Beacon 1963-7 : A System Design Ahead of its Time	24
<i>Michael Knight</i>	
The Ghost of KDF9	34
<i>Brian Wichmann</i>	
Letters to the Editor	36
Obituary: David Caminer	38
<i>Frank Land</i>	
Obituary: Professor Hilary Kahn	39
<i>Dik Leatherdale</i>	
Forthcoming Events	39

Greetings!

Dik Leatherdale

This is my first issue as editor of *Resurrection*. Nicholas Enticknap has guided *Resurrection* since it first began 18 years ago, establishing an enviable (and for me, daunting) standard. If I can follow an act like that, then I shall be doing exceptionally well. Or rather, we shall be doing well. After all, it's not the editor who writes the articles in *Resurrection* it's you, the readers. I know that you will continue to support *Resurrection* with as much enthusiasm as you have over the last 18 years.

An innovation in this edition is the first of a series of profiles of significant figures in the world of computing. CCS members will, I feel sure, be familiar with the achievements of the three great pioneers – Turing, Kilburn and Wilkes. But there are many important people, of whom we have all heard, but whose work is not as well known as it should be. In this edition, we make a start with Professor Christopher Strachey, a seminal figure who is, perhaps, not now as familiar to us as he deserves to be.

Chairman's Report for 2007-08

David Hartley

At the Society's AGM, held at the Science Museum on 15 May 2008, I was able to report much achievement and progress during the year passed.

The CCS held the usual full programme of meetings and lectures both in London and Manchester, two of which stand out as worthy of special mention. Under the leadership of our immediate past Chairman, Dr Roger Johnson, a three-day seminar in July 2007 was organised to celebrate the 50th anniversary of the founding of the BCS. The first day was held at Bletchley Park and featured Colossus and the recently completed Turing Bombe. The circus then moved to London where, in two days, we covered the rest of the British story including the Manchester and Cambridge achievements, the early developments of Ferranti and Elliotts, as well as milestones in the development of a professional society; and much else besides. Then in March 2008, pioneers from the BBC Computer Literacy project and from Acorn Computers reviewed the legacy of the BBC Micro. Although barely 25

years old, it is a tale that tells a remarkable story in the history of computing.

CCS projects continued to restore or rebuild old equipment, preserve artefacts (software as well as hardware) and record past events for posterity. Most notably the Turing Bombe re-build, led by the indefatigable John Harper and his large team of volunteers, was completed in time for the BCS@50 meeting, and for a switch-on ceremony by the BCS patron, HRH the Duke of Kent. We are proud that the project was formally recognised by the BCS with the award of an Honorary Fellowship to John Harper.

Another re-build project, the Manchester Small-Scale Experimental Machine, or “The Baby” as it is affectionately known, which was completed 10 years previously, found a new home in the Manchester Museum of Science and Industry. It is now placed at the very entrance to the museum, so that no one can miss this achievement of the pioneers of 1948 nor that of Chris Burton and his team in 1998,

The opportunities for the display of artefacts, as well as for further restoration work, were enhanced during the year by the foundation at Bletchley Park of the National Museum of Computing. The success of the Science Museum project in constructing Babbage’s Difference Engine No. 2 was continued when a second copy was shipped to the Computer History Museum in California. Your chairman was able to attend the events associated with this in May 2008.

Finally, two matters concerned with communications. The CCS website was re-vamped at the start of the year with the aid of a new webmaster (Kevin Murrell) and a new editor (Alan Thomson). Thanks to their efforts, there is no doubt that we are serving our members and others better than hitherto. We announced the retirement of, and paid tribute to, Nicholas Enticknap after 18 years and 42 issues as editor of *Resurrection*. As I said in issue 42, the Committee was working hard to find a successor. We did not have to look far, because very quickly, Dik Leatherdale, an active member of the CCS, volunteered for the role. The Committee was equally quick to take up the offer, and with this issue 43 you can see the wisdom of their decision.

News Round-Up

We have been contacted by the newly-established Centre for Computing History in Haverhill near Cambridge. The Centre specialises in post-1975 computer history but the preceding years are not overlooked. Members who have an interest in “modern” computing are encouraged to visit and/or to explore their website at www.computinghistory.org.uk.

101010101

But as one door opens, another closes. The Museum of Computing in Swindon (www.museum-of-computing.org.uk) has lost its premises and has put its exhibits into store for the time being. It is optimistic that a new venue will be found.

101010101

Alarming news from Bletchley Park in *The Register* (www.theregister.co.uk). In May, a story appeared which suggested that the entire project was in danger of financial meltdown. Your intrepid editor made some enquiries via Pete Chilvers and was reassured to find that, although the situation is finely balanced, careful management was keeping matters under control. Some restoration of Turing’s Hut 8 and urgent repairs to the mansion roof have, nevertheless, been possible.

At this point the story in *The Register* disappeared.

Nevertheless, Bletchley Park does need to raise substantial funds for restoring and maintaining the fabric and is negotiating with various possible sources. A well-wisher has set up an online e-petition to the Prime Minister at petitions.pm.gov.uk/BletchleyPark which members are encouraged to sign. Part of the original *Register* story can be found there too.

As *Resurrection* went to press, the Bletchley Park petition ranks 22nd of over 6,000 active petitions within Downing Street’s e-petition system with over 5,800 signatures. It would be churlish, however, to draw attention to the fact that the e-petition system has been in beta test since November 2006.

In July it emerged that discussions with the Heritage Lottery Fund have been encouraging, which may allow some optimism that money from this source might be forthcoming in the fullness of time. Matching funds, however, will still need to be raised to which end a new fund raising campaign will be launched in the autumn.

101010101

Manchester Digital 60 Celebrations

Chris Burton

The 60th anniversary of the first successful running of an electronically stored program on the University of Manchester (UoM) Small-Scale Experimental Machine has just been celebrated on a full day of events at the University and at the Museum of Science and Industry (MOSI). We actually celebrated on Friday 20 June, which was more convenient than the correct anniversary on Saturday 21 June.

The University PR people had done a good job in the lead-up to the day, with press coverage of the anniversary itself, as well as releases about “newly-discovered” items which are quite familiar to us in the CCS. One item, which generated a lot of interest, was the recording of music fragments generated by the Ferranti Mark 1. The disc was donated to the CCS by the late Frank Cooper and is described in *Resurrection 12*. This drew the attention of the BBC Today programme, and your correspondent was briefly interviewed by Jim Naughtie on Wednesday 18 June. The other “discovery” was the well-known photograph of the prototype Mark 1 which appeared in *The Illustrated London News* in June 1949, and which was a composite of Alec Robinson’s photos of 15 December 1948. A spot on BBC Breakfast TV on the Friday also helped raise the profile.

The School of Computer Science (CS) at UoM had prepared a Schools Digital 60 Day for schoolchildren. Almost 500 schoolchildren attended in a big new auditorium. During the morning they were treated to a video link-up from MOSI, where Geoff Tootill and your correspondent were interviewed by charismatic Professor Brian Cox.



After the video link, there was a media scramble and photo opportunity at MOSI, with the other pioneers Dr Alec Robinson, Professor Dai Edwards and Professor Tommy Thomas represented by his grandson Ben. MOSI has made an excellent job of

enhancing the SSEM site with quality signage and safety barriers.

The CCS (NW) group organised a Celebratory Seminar, chaired by Tom Hinchliffe. We had the pleasure of hearing a talk by Bill Olle - "Eight Significant Events in 60 Years of Computing", followed by Professor Simon Lavington - "Where did it all go? - Tracing the Influence of Manchester Computer Designs, 1948 - 1974".

An evening reception hosted by CS then followed, with poster displays, and the first outing of an amazing full-size cardboard cut-out model of the SSEM. This is an assemblage of full-scale photographs of the racks on a pop-up display frame. The centre rack has a couple of embedded processors emulating the SSEM, with appropriate display and control switches. It is almost like operating the real thing (except it isn't warm enough!), and can be folded-up to take to exhibitions.

Before the start of the ensuing prestigious Kilburn Lecture, a Manchester Literary and Philosophical Society event, there was a medal award ceremony for the four pioneers. The University of Manchester awarded each of them a University Medal of Honour - the highest non-degree award. These medals were presented by the Head of the School of Computer Science, Professor Chris Taylor. Each pioneer was then presented with a BCS@50 medal which had been struck in 2007 to commemorate the 50th anniversary of the British Computer Society. The four medals had been engraved with "The Baby at 60" and were presented by our Chairman Dr David Hartley. After the presentations, your correspondent was delighted to be himself honoured with a BCS@50 medal, commemorating ten years of the SSEM replica being shown to the public.

[The CCS adds its hearty congratulations to the four pioneers and to Chris - Ed.]

The Kilburn Lecture "The March of the Microchip" was given by Professor Steve Furber. This excellent lecture can be found on the website below. The lecturer highlighted the extraordinary development of electronic technology over the last 60 years, particularly as applied to computing devices. Provocatively, he concluded that technology today is capable of being used to start mimicking brain functions.

The website www.digital60.org is intended to be a permanent record of the achievements of the University in the development of computing.

101010101

July also saw a rare public exhibition of the preserved ICT 1301 at a classic car rally in Kent. See www.ict1301.co.uk.

101010101

Babbage's Difference Engine No. 2 Arrives in California

Doron Swade

Babbage's Difference Engine No. 2 is on display at the Computer History Museum (CHM), Mountain View, California. The Engine, a duplicate of the one on display at the Science Museum in London, arrived on 9 April and has been something of a sensation. The machine was commissioned by Nathan Myhrvold, formerly Chief Technical Officer of Microsoft, for his private collection. Myhrvold's Engine, like the original in London, was built by the Science Museum and is on loan from Myhrvold for one year to the Computer History Museum where it will be demonstrated near-daily by volunteers.

The Engine is the centre-piece of a temporary Babbage exhibition which is part of the entrance feature of the Museum. The arrival and unpacking of the five-tonne Engine created a stir in the Valley and was well covered by press, blog, TV and radio. The exhibit and accompanying website [www.computerhistory.org/babbage] were curated by Doron Swade, Director of the Babbage Project and Guest Curator of the exhibit.

The public launch of the exhibit on Saturday 10 May was a Babbage fest and it drew capacity crowds. Some 1800 people attended. In the spirit of the occasion, some came in Victorian costume. The day-long programme included the official opening of the exhibit, demonstration of the Engine, a lecture by Doron Swade, two screenings of the documentary 'To Dream Tomorrow' (an authoritative treatment of Ada Lovelace and her collaboration with Charles Babbage), and free popcorn.

In the run-up to the exhibition the Computer History Museum featured a platform session 'An evening with Nathan Myhrvold and Doron Swade discussing Babbage's Difference Engine'. The event, held on 1 May, was oversubscribed, and the auditorium, which seats 450, was at capacity. Myhrvold and Swade gave presentations followed by a question-answer session mediated by Len Shustek, Board Chairman and acting Executive Director of the Museum. Myhrvold expressed delight at finally receiving his Engine and confessed that, over the years, he had had private doubts as to whether he would ever see its delivery. CHM was delighted to welcome David Hartley, CCS Chairman, to this event, as well as to a private VIP Babbage event at CHM on 29 April.

Myhrvold's Engine was built by Richard Horton and Chris Ward at the Science Museum, and the construction, which started in 2002, took six years. The Engine consists of 8,000 parts, weighs an estimated five tonnes and measures 11 feet long and seven feet high. Like its twin at the

Science Museum, the machine includes the printing and stereotyping apparatus, an integral part of the original designs.

Little was known about how the Engine would fare in trans-Atlantic transit by air. In the event it fared badly. The machine was built to Victorian instrument-makers standards and was in impeccable exhibition condition when shipped from the Science Museum. Its condition on arrival was dismaying. Steel and bronze parts had corroded, a response, it is thought, to extreme low temperatures at high altitudes – this despite precautionary spraying with protective silicone. There was worse. The mechanism had been shaken to bits – literally. Machine screws had shaken loose and fallen out and at least one tapered pin was found loose in the oil pans below the mechanism. The mounting plates that exactly position the eight columns of figure wheels had all loosened and the precise and painstaking adjustments that had taken months were lost. The Engine had also clearly been dropped in transit.

The shipment was met at CHM on arrival by Horton and Fulcher from the Science Museum, and Doron Swade, all of whom arrived in California the day before. Horton and Fulcher were scheduled to spend two weeks setting up and testing the machine, and training the local docent demonstrators. With only two weeks to make good the effects of transit damage, there was no question of dismantling the machine and redoing the detailed adjustments. Under Horton's supervision three CHM volunteers set to with Brasso and polishing cloths to remove as much corrosion as possible while Horton and Fulcher set about restoring many of the fine adjustments from a largely unknown starting point. Horton and Fulcher were assisted by Tim Robinson, an ex-Brit CHM volunteer, who heads the CHM Babbage support programme. Robinson is no stranger to Babbage's designs: he has already built a working Babbage Difference Engine No. 1, Difference Engine No. 2, and parts of the Analytical Engine in Meccano [www.meccano.us].

By the time the Science Museum engineers left at the end of their fortnight, the Engine could be demonstrated reasonably reliably. There were residual problems of uncertain severity and a few breakages, but workarounds were found to allow the Engine to be operated at all the scheduled public events, which it did with some success.

The delivery of a working Engine to California completes an undertaking made in 1997 to Myhrvold that the Science Museum would build him a Babbage Engine as part of a benefaction arrangement that funded the construction, completed in 2002, of the Science Museum's Engine. The duplicate Engine is destined for Myhrvold's living room in Seattle after its year at the CHM.

Society Activity

Pegasus Working Party

Len Hewitt and Peter Holland

In the last issue we reported that Pegasus had problems with its cooling system and with the mains power system. A new cooling compressor was fitted on the 14 June. All the problems have now been resolved and Pegasus appears not to have suffered any ill effects from an incorrectly rotating alternator and confused switch-on sequences. We were all extremely relieved when the machine ran faultlessly for five hours or so on the 18 June. That was the first long untroubled run since last September. Our “In Steam” days are every other Wednesday from 2 July switched on from 11:00 to 15:00.

Contact Len Hewitt at leonard.hewitt@ntlworld.com.

Bombe Rebuild Project

John Harper

Now that our basic machine is complete and available for display and demonstrations, this is an ideal time for reflection. I would like to leave my technical report to next time and to use this opportunity to say a massive **thank you** to all those people and companies who have helped us succeed in completing this mammoth rebuild task over a 12 year period.

Sadly some of those who worked on the original Bombes, either in manufacturing or maintenance and who helped us as consultants in the early days, are no longer with us. Also we have lost touch with a few people over the years and I cannot therefore thank them personally. One or two companies have also gone out of business over the years, but perhaps this report will get to the individuals who helped us.

As was reported in the last issue, I have been awarded an Honorary Fellowship of the British Computer Society, but I hasten to add that this award is really for the whole 60 or so people of our team who made this possible. Clearly I could not have even started without the help of all those involved. It says a great deal about those early team members and the companies and supporters that they were so willing to help when we had only a plan and little else.

The full list of team members and supporters can be found at www.jharper.demon.co.uk/bombel.htm.

Last autumn at the annual reunion of the Bletchley Park veterans we had another 'Switch On' ceremony where many WRNS operators and other similarly involved during WWII were present. They all appeared to be very pleased with what we had achieved, bearing in mind that the main reason for carrying out this rebuild was to pay tribute to all those whose story could not be told until about 30 years after the war had ended.

Our most recent event was when the Fellowship was formally presented at BCS London HQ on April 23. I would wish to thank the CCS and the BCS for this. Many of our team members and their partners were able to attend, but by no means all. Those who were not able to be present are also to be thanked and this I do personally but also on behalf of the CCS and BCS. The photograph below is of those members of the team who were able to attend plus CCS and BCS officers.



Left to Right - Simon Harper , Hester Walden , Vi Maile* , Tony Walden* , Rachel Burnett (BCS President) , David Hartley (Chairman CCS) , Janet Harper , Chris Tarry* , John Harper* , Roger Johnson (past Chairman CCS) , Ian Walker* , Mary Hillyard* , Mike Hillyard* , John Borthwick* , Di Borthwick
* Team member*


Pioneer Profiles – Christopher Strachey

David Barron

Author's Note.

This is not a fully-referenced scholarly paper. Rather, it is an affectionate tribute to a former colleague and friend and is based mostly on memories of conversations, whether at High Table in Cambridge, in the Laboratory, or in the rural seclusion of the Strachey family home in Sussex.

Anyone who recognises the name 'Strachey' will no doubt associate it with the literary family who were at the heart of the Bloomsbury Group (think 'Eminent Victorians' by Lytton Strachey – Christopher's uncle). And those few computer scientists who recognise 'Christopher Strachey' will probably associate the name with his work on formal semantics of programming language at Oxford, in collaboration with Dana Scott. But Christopher was far more than a theorist: he was always the programmer's programmer, and also played a leading part in the founding of the British computer industry, as a logical designer.



Although the Stracheys were mainly a literary clan, there was a mathematical streak in the family – Christopher's father Oliver was engaged in decryption in both world wars. Christopher followed this streak, and went up to King's College Cambridge in 1935 to read Mathematics. After various vicissitudes, he graduated in Natural Sciences and took up a job as a research physicist with STC (Standard Telephone and Cables Ltd.). At the end of the war, he left STC and took up teaching, eventually becoming a mathematics master at Harrow School in 1949.

During his time with STC he had been involved in numerical solution of differential equations using a Differential Analyser (an analogue computer). Whilst at Harrow, he was introduced to Mike Woodger who told him about Turing's 'Pilot Ace' computer at NPL. Christopher wrote a program to play draughts on the machine, but the Pilot Ace wasn't really up to the job. Hearing about the Manchester Mark 1, he wrote to Turing, his contemporary at King's, asked for details of the

instruction sets and completed a program that not only played draughts, but also played “God Save the King” on completion.

In 1949, Lord Halsbury had persuaded the Government to set up the National Research Development Corporation (NRDC) under his leadership, the intention being to commercialise British scientific ability. Halsbury was particularly seized by the potential of the then new computers, and persuaded Strachey to join NRDC. As well as doing the programming for a simulation of the proposed St Lawrence Seaway, he took a major role in the development of the Elliot 401 and Ferranti Pegasus computers, being responsible for the logical design of the Pegasus, a workhorse to replace the Ferranti Mark 1 (based on the Manchester Mark 1). In 1959, he left NRDC and set up shop as the country’s first freelance computer consultant. In this capacity he had substantial input into the design of EMI’s EMIDEC 1100 and 2400 computers.

In 1962, Strachey, whilst continuing as a consultant, was persuaded by Maurice Wilkes to join the team at Cambridge developing a cut-down version of the Manchester Atlas supercomputer (called Titan in Cambridge, but marketed by Ferranti/ICL as Atlas 2). Strachey’s brief was to develop a new programming language for the machine, working with myself and David Hartley. The language was based on Algol 60, and was initially called CPL (Cambridge Programming Language). Later, we joined forces with a team at the University of London Institute of Computer Science, and it became ‘Combined Programming Language’. For myself, I am content for it to be remembered as Christopher’s Private Language.

CPL had many innovative features: some of these were just, in Christopher’s phrase, “syntactic sugar”, but a major contribution was the clarification of the concept of L-values and R-values, which can be seen in C and all subsequent languages. (CPL begat BCPL, which begat B and then C and C++: another example of the pervasive influence of Christopher at the time). The development of CPL also provides another instance of Christopher as a programmer’s programmer. He decided that we needed a macro generator to assist in developing the CPL compiler, and – over a weekend in his Sussex home – produced the General Purpose Macrogenerator, GPM. This was an incredibly elegant string-substitution language – which could also, as he demonstrated in a tour-de-force of programming, be used to compute factorials.

As the CPL project proceeded, he became more and more interested in the formal semantics of programming languages. Delivery of the

compiler fell more and more behind schedule. As a result, in 1965 he accepted an offer of a post at Oxford, as Head of the Programming Research Group (PRG), an offshoot of the Computing Laboratory. Here, in collaboration with Dana Scott, he developed his theory of denotational semantics and was eventually recognised by the award of a Personal Chair. But even whilst he was engaged in this theoretical work, the demon programmer survived. The PRG was given funds to buy a Modular One computer. Christopher decided from the start that the group would do their programming using an interpreter to simulate a stack machine. This was to be embedded in an operating system (OS/6) based on high level concepts. In the months between the placing of the order for the Modular One and its delivery, Christopher and his assistant Joe Stoy wrote the system in a CPL-like language, and developed a cross-compiler on the University's KDF9 mainframe to compile the code into Modular One assembler. The machine eventually arrived, and the engineers installed it. The cross-compiled code was loaded, and pretty well worked out of the box. A night of tweaking followed, and in the morning the PRG personnel appeared to find a working system. Strachey locked away the Modular One manuals: "I'm the only one who knows the instruction code" he joked. His premature death in 1975 deprived the country of one of its greatest and most prolific computer scientists. (He always denied the existence of Computer Science, but he will be remembered as one of the subject's founding fathers.)

There's so much more that I could say about this remarkable man: perhaps the Editor will allow me another thousand words in a forthcoming issue. Let me finish with a quotation that says it all:

"It has long been my personal view that the separation of practical and theoretical work is artificial and injurious. Much of the practical work done in computing, both in software and in hardware design, is unsound and clumsy because the people who do it have not any clear understanding of the fundamental design principles of their work. Most of the abstract mathematical and theoretical work is sterile because it has no point of contact with real computing."

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*

Computers at Metrovick : the MV 950 & AEI 1010

Ron Foulkes

Ron Foulkes gives us an excellent description of the early computers at Metropolitan Vickers, one of the least known British computer manufacturers, and his role in their development.

I will start by telling you how I got involved with computers and, tied in to this, why Metrovick/AEI got involved. I will go on to describe the development of the Metrovick 950 followed by the AEI 1010. Finally, I will describe a particular 1010 installation.

A point of clarification: up to 1957 although the Metropolitan Vickers Electrical Co (commonly referred to as Metrovick, Metros or MV) was part of Associated Electrical Industries (AEI) it operated under its own name, as did British Thomson-Houston (BTH), the other major part of AEI. In 1957, the Metrovick and BTH names were dropped. This is why I start by referring to Metrovick and then later to AEI.

I have had the help of my original design notebook and a copy of an article I wrote for the MV Gazette in 1957. The notebook covered the period mid 1954 to mid 1958 and recorded my thought processes and calculations during the design of both the 950 and 1010. I have also borrowed some sales literature, software and maintenance manuals for the 1010 from John Hardman, who was a senior installation and commissioning engineer at AEI on 1010 systems from the late 1950s to the mid 1960s. The commissioning and maintenance engineers were the unsung heroes of those early days.

Starting with how I got involved: I grew up in Llandudno N. Wales where most people made a living out of holiday-makers or sheep farming. I was the only one of my school contemporaries who went on to study engineering. I had an uncle who worked in the Research Dept of Metrovick and he convinced my father that I should go to Metrovick because engineering was interesting and fun.

So I went to Metrovick in 1946 and did a one year probationary college apprenticeship, followed by a three year electrical engineering degree at Manchester University and a final year college apprenticeship back at Metrovick.

My university degree course was virtually all related to 50 cycles/sec (Hz) although I was aware that something to do with computers was going on in a room I occasionally passed.

When I returned to Metrovick I moved between departments – steam turbines, alternators, transformers, switchgear etc. – and accepted a job in switchgear engineering to start in Jan 1952. I should have finished my apprenticeship there, but an administrative error sent me to the Radio Engineering Dept. and I was persuaded to stay.

There were three people in the department who were a major influence on me, and therefore on my work, after I finished my apprenticeship.

Glyn Jones was a senior engineer working on display systems. During my time with him we worked on a large air defence ground radar system, developing the first synthetic display which is the basis of today's air traffic control systems. Instead of a dot on the screen, the operator saw a packet of information about the target.

Jack McQueen was a brilliant innovative designer whose hobby, apart from his work, was building HiFi amplifiers. We reckoned he spent more time looking at the output on an oscilloscope than listening to them! No one had ever seen what a radar pulse looked like – it was too short to be seen on the current oscilloscopes, so Jack was asked to solve the problem. He did so by developing the first sampling oscilloscope.

In both cases, my contribution was mainly wielding a soldering iron and performing “suck it and see” experiments, but it gave me my grounding in electronic system design.

Dr L. W. Brown was Chief Engineer of the Radio Engineering Dept whose business, when I joined, was confined to military radar and control systems. In late 1954 I was called to his office and asked to lead a small team to develop a computer for solving mathematical and scientific problems. He was always approachable and I particularly recall two lessons he instilled into me:

- Never ask for permission unless you know the answer will be yes.
- Permission is always harder to get than forgiveness.

If I asked for a decision, I seldom got one from him, but after I had left his office, I always knew what it should be.

These three people will never be directly associated with early computer development, but their influence on me was crucial to my subsequent contribution.

I think the background to Metrovick starting to design computers (which I did not know at the time) was –

- Dr Brown worked in the Telecommunications Research Establishment (TRE) during the war and was familiar with what his contemporaries were doing in the computer field.
- Metrovick was run by engineers for the benefit of engineers – well it seemed like that!
- Metrovick engineers were some of the first people to make use of the Mark 1 machine at Manchester University for such applications as stability of transmission networks, radar aerial design, turbine blade design and critical speed calculations.

Soon after Lord Chandos had returned as Chairman of AEI in 1954 he told the shareholders that the future expansion of AEI would be equipment for nuclear power stations, railways and automation.

So I now believe that Metrovick/AEI knew what it was doing. It understood that computers were going to play an essential part in its traditional business and it was making sure it was on the computer train, although it was not sure where the train was going.

The Metrovick 950



The Metrovick 950

The late Professor Dick Grimdsdale described his work on developing a transistor computer in *Resurrection 13*. Dick did a degree in Electrical Engineering at Manchester at the same time as I. I did not know him well as we had different interests. Dick worked hard and got a first. I was distracted by girls

and sport and got a second. It was not surprising that he stayed on to do research and had a very successful academic career.

There is no doubt that his work was the inspiration for the Metrovick 950 and I don't want to minimise that in any way by what I am going to say next. Nor the work of G.B.B. Chaplin, who had been in our year and who was investigating point contact transistor circuits at the University.

I visited Dick who was working on his transistor computer at Manchester University. I also visited Dr A.D. Booth at Birkbeck College and I think we bought a drum for our first experiments, which his father made. It was an aluminium wheel about seven inches in diameter and 2½ inches high. It was not magnetic until we sprayed it with iron oxide.

Discussing his design with Dick was a great starting point. but I do not recall receiving any detailed design documentation from Dick and my design notes show that we deviated quite a bit from his design at the very outset. For example, the 950 word length was 36 digits compared to 48 and the clock frequency was 57KHz compared to 125KHz. The lower frequency made a significant contribution to reliability.

A two address instruction was used, but instead of the second address being the destination of the result, it was the address of the next instruction. This format had been used by Dick and enabled the effective access time of the drum to be reduced.

The registers were regenerative tracks on the drum as used by Dick, but the lower frequency increased the distance between the read and write heads, allowing wider tolerances and better reliability.

One novel feature of the 950 was that instructions could be read directly into the instruction register rather than via the accumulator, a suggestion by Gerry Proctor, one of our first programmers.

Addition and subtraction took three milliseconds. Multiplication, using eight adders in parallel, took eight milliseconds. Division by subroutine took 750 milliseconds. This compares with one second quoted by Dick for his machine which had a higher clock frequency – 125KHz compared to 57KHz. I don't know why the 950 was faster.

I was amused to read in the article I wrote for the MV Gazette in May 1957 that “an error will occur if an arithmetic operation gives an answer outside the range of the computer. A programmer should not allow this to occur but if it does, a lamp on the control desk will light up!”. Perhaps being an engineer at Metrovick made one a little too arrogant!

I was assisted on the initial hardware development by Geoff Atkinson and Bill Loughhead. John Bailey joined the team later. The

input/output peripherals were the responsibility of John Gladman. The first programmers were Ken Evans and Gerry Proctor.

The drum was a key component. There was a high level of mechanical design and manufacturing skill available to us. I believe John Smethurst was the mechanical engineer who designed the drum – aluminium – eight inches high, five inches diameter. It used angular bearings under pressure and was balanced and machined to a tolerance of 0.0002 inch. Heads were punched from 0.015 inch mumetal lamination with a 0.001 inch gap defined by aluminium foil. It was surprisingly reliable.

The printed circuit boards (PCBs) were made in-house. We used point contact transistors from STC in the first production model. It used approx 230 transistors and 90 PCBs. Mullard OC71 junction transistors became available and these were used in the later six production models.

Reliability was adequate, but I cannot quote any figures. The principal problem was the PCB connections which relied on the socket contacts cutting into the surface of the copper on the PCB.

One computer was supplied to Queens University Belfast, one was in the AEI computer centre in Sale and the other four were in various departments of Metrovick.

The AEI 1010

In 1956/7 we began to think about the successor to the 950. The availability of the transistor was the basis of the design of the 950. I don't think the ferrite core is given the credit it deserves in the progress of computer design. Its availability and its properties were essential to a number of key design features of the 1010.

By this time the growing market for computers was seen to be data processing and this guided our thinking.

- A microprogrammed central processor, with a 44 digit word length and floating point arithmetic.
- A 4096 word fast access ferrite core store.
- An 8192 word magnetic drum store.
- Up to 31 peripherals connected via a common highway.
- Parallel programming (time sharing).

I am not concerned about claiming firsts for our work but reading "Timesharing History" by John Deane in *Resurrection 31* and Alan Thomson's subsequent, "Timesharing History: the UK Story" in *Resurrection 32*, it does appear that we were up there somewhere!

No longer restrained by drum access time, we used a single address format.

It was envisaged that a wide range of peripherals would need to be used. Magnetic drums and additional core stores were treated as peripherals.

Initial circuit design work used OC71 transistors, as used in the Metrovick 950. As the range of available transistors grew, these were used as the design progressed and incorporated in the production design. These included GET 102, GET 872, GET 882, GET 885, GT50, TK36C and 2N598. Thermionic valves were only used to generate the master timing pulses at low impedance.

The core store had a capacity of 4096 44 digit words plus four parity bits with an access time of 3.5 μ sec and a read/write cycle time of 8.5 μ sec. There was also a small core store containing eight 13 bit B registers.

Loose ferrite toroids about 3mm in diameter had become available. Girls in our factory knitted them into arrays for circuit development purposes. Later, in our production units, we used arrays manufactured by Plessey.

The main arithmetic unit performed 44 bit parallel arithmetic and a second arithmetic unit performed 13 bit parallel arithmetic.

There were two main registers each 44 bit. For multiplication and division they were combined to hold a double length number. A further 44 bit register was used in multiplication.

The control unit extracted instructions from the core store as specified by the 13 bit long address in the control register. There were finally about 86 instructions the programmer could use. As microprogramming was used, these were all built up from a basic set of operations. and therefore. the logic of the control unit was able to be designed before the final instruction code was settled.

The data scanner controlled all the transfers of data to and from the peripheral equipment. Up to 31 peripheral devices could be attached to a common highway. Data was transferred in blocks of 32 words.

The original concept was to have a buffer consisting of a one word staticiser and a ferrite core store with a capacity of 32 44 bit words for each peripheral, together with appropriate control logic. A 32 word block would be transferred serially along the data highway at the data rate of the peripheral. Each word would be staticised and transferred to the core

store. The data scanner continuously scanned the state of buffers – empty, full or in use – and initiated the high speed transfer of data between the buffer and the central processor store – outgoing data to empty buffers and incoming data from full buffers.

Later in the design process, the core store was dispensed with. The complete cycle of identifying a full one word buffer and transferring the data to the central processor store was 26 μ sec but, while that data transfer was taking place, the search for the next full buffer continued and the net time was 13½ μ sec. The drum was the peripheral with the fastest data rate at 600 μ sec per word, so a transfer only occupied 2% of the central processor time.

AEI 1010 Peripherals

The following devices were envisaged:

AEI type 972 magnetic drum 8192 words	Decca type 3000 magnetic tape deck
AEI 1000 ch/sec paper tape reader	Ampex FR-300 magnetic tape deck
AEI 250 ch/sec paper tape reader	Bull tabulator
Elliot card reader	ICT Samastronic lineprinter 300 lpm
Elliott card punch	Rank Xeronic lineprinter 3000 lpm
Creed Model 25 reperforator	Enquiry typewriters
	Creed Type 3000 paper tape punch

Parallel programming

I have seen, in the literature, a lot of discussion and claims as to what in those days was often called “time sharing”. In the 1010 there were two levels of time sharing the central processor.

1. Transfers to and from peripherals.

This only involved time sharing access to the main core store between the main registers/arithmetic units and the data scanner.

2. Time sharing the central processor between application programs.

We referred to this as parallel programming. If, for example, the central processor was held up waiting for data from a magnetic tape deck, it could switch to another program. Another example would be where a long running application program had to be interrupted to

allow some urgent other work to be done. It seemed to us an obvious requirement once magnetic tape decks were to be used.



Three systems were sold to the National Coal Board in the late 1950s. Two were installed at Sighthill, Edinburgh and one at Lowton, Lancs – mainly used for payroll.

All six systems were taken out of service soon after the GEC takeover of AEI in 1967.

AEI 1010 at NCB Edinburgh

RAF Hendon

In the early 1960s, a dual processor system was purchased by the Air Ministry to provide a global stock control system for the Royal Air Force. I will describe this in more detail.

At that time, the RAF had 150 stations/depots, home and overseas, holding and requiring spares for maintenance which were being consumed at a rate of £40M a year. The purpose was to ensure that the right equipment was in the right place, at the right time.

The installation at RAF Hendon would hold records of stock holdings, outstanding orders etc., updated daily on magnetic tape.

Each station/depot reported all local transactions by teleprinter daily. This resulted in a paper tape record being produced at Hendon and then used for the daily update of the central records. Urgent priority demands were transmitted immediately they arose and were dealt with within the hour at Hendon. This was an application of “parallel programming”. The daily update program ran for up to six or eight hours a day and was not stopped by the operator to allow priority demands to be dealt with.

The 1010 installation was run continuously on a three shift basis.

It was said that the capital investment in the system was £2.5M and resulted in savings of £1M a year.

The computer system comprised:

- 2 AEI 1010 computers
- 3 AEI magnetic drums (8192 Words each)
- 1 AEI ancillary core store (8192 words)
- 24 Ampex half inch 45 KHz magnetic tape decks
- 4 AEI high speed paper tape readers (1000 ch/sec)
- 4 Creed paper tape punches (300 ch/sec)
- 1 Xeronic line printer (3000 lines/min)
- 1 Anelex line printer (1000 lines/min)
- 1 Elliott card reader (370 cards/min)



AEI 1010s at RAF Hendon

Each 1010 had its own highway and each peripheral could be connected to one or the other by plug and socket.

With hindsight, I think the Hendon project stretched the capabilities of the 1010 too far. In particular, connecting so many peripherals meant long highways, which gave

timing and loading problems. The Air Ministry, not unreasonably, demanded rigorous acceptance tests and it took almost four years from assembling the system in the factory to it being fully commissioned and operational. The site was finally completely operational in 1966. This took up a huge amount of the design and commissioning team's time.

Nevertheless, the final outcome was claimed by the customer to be one of the largest stock control systems in Europe. It operated 24 hours a day for the 10 years up to 1976. It was estimated that about 20 miles of punched paper tape was read into the system each day and a further four miles of punched tape was output daily for transmission to RAF stations and depots world wide.

959 and 1040

Two other machines were developed.

The 959 was designed to be installed in hostile environments. It was a collaborative project between AEI and Stewarts and Lloyds. The one

and only installation was at Stewarts and Lloyds in Corby, where it was adjacent to and controlled a flying shear at the exit of a rolling mill producing steel bars. The application program took the order book and optimised the operation of the shear to minimise wastage.

The computer was mounted in a box the size of a coffin and was water cooled!

The 1040 was a 1010 with input/output scanners supplied by AEI Harlow. One was installed at Oldbury nuclear power station to provide a comprehensive alarm analysis system for the power station. It was commissioned in 1965.

The End of the 1010

In 1965 AEI decided to concentrate on the industrial process control market and the data processing activities were closed. A decision had been taken in 1964 to negotiate an agreement with GE of America to build their range of GEPAC process control computers under licence. Several systems were built in Leicester and installed under the name AEI CONPAC.

In 1967 GEC took over AEI and then, in 1968, absorbed English Electric, which had itself divested itself of its data processing business earlier in 1968 into ICL. Thus, GEC did not ever have or inherit a data processing business. It did, however, end up with a multitude of other computer activities in AEI, English Electric, Marconi, and Elliott Automation – in addition to its own – but that is another story.

Editor's note: This is an edited transcript of a talk given by the author at the Manchester Museum of Science and Industry on 12 February 2008. Some additional photographs of the 1010 can be found in Resurrection 38. Contact Ron Foulkes at ronfoulkes@dsl.pipex.com.

CCS Web Site Information

The Society has its own Web site, which is now 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 this latter URL is case-sensitive.

Beacon 1963-7: A System Design Ahead of its Time?

Michael Knight

Michael Knight recalls Beacon – the pioneering reservation systems developed by BEA and Univac. Despite achieving so much with so little, Beacon did not survive the BOAC merger, but its ghostly legacy lingered on for a while in unexpected places.

Beacon was British European Airways' computer online network, initially developed in 1963-7, to provide a full-scale passenger reservations service. Subsequently, the hardware was upgraded, and further applications added on an integrated basis. Later, following the merger with BOAC to form British Airways, these services were progressively taken over by Beacon's old rival system Boadicea. The memory of Beacon thus began a long evaporation.

The purpose of this memoir, 40 years on, is to record something of the achievement of the Beacon project team on what was, for a time, Europe's largest and most ambitious business transaction-processing (TP) system. TP is the norm now, taken for granted in systems of any size, but back then was rare and leading-edge, not even called TP, but 'real-time'.

The first part outlines the task, and the resources deployed to implement it: hardware, software, time and people. The second part sets out to justify this paper's bold title, by examples rather than by a systematic description of the design. Some of these examples were subsequently reflected in standard operating system facilities for TP applications; some were not. In the latter case, this has been due to the awesomely rapid developments in computing hardware capabilities, which have allowed the evolution of increasingly function-rich and hardware-profligate operating systems and other software tools. In the early 1960s, before integrated circuits and the ubiquitous byte, when ferrite core memories were hand-knitted, programmers were far from cheap. But hardware was unimaginably expensive by today's standards, and limited in its capability, and there were definite limits to what could be had in one box, at any price.

The Task

In the mid-1960s, BEA carried over 7,000,000 passengers a year, and was 5th largest 'carrier' in the world. (BOAC carried over 1,000,000 a year and was ranked 35th largest). By today's airline standards, these numbers are very modest, but they presented very demanding data

processing challenges. The real-time solutions to these challenges were the ancestors of today's all-pervading call-centres.

Airline reservation systems involve up-to-the-minute updating of numerical records of flight occupancy – or, more accurately, of each flight/class/segment/date, in BEA's case for 11 months ahead. This was known as 'inventory'. Related to that, alphanumeric passenger record details need to be formed and maintained, a record being for one or more named passengers travelling together on an itinerary of one or more flights. This was known as a PNR – passenger name record. Linking PNRs to relevant flight records were PNIs – Passenger Name Indices.

Bookings, amendments, etc. were typically made by telephone, by airline or travel agents, or by travellers. An added complication was that airlines made bookings with each other, on behalf of travellers, and needed to be kept aware of the booking status of many of each others' flights, as well as traveller details. The medium for these 'interline' bookings and status updates was a jointly-owned international airline message network, SITA, carrying mainly manually-originated teleprinter messages on 75bps asynchronous circuits, the messages conforming (or more often not) to a common protocol called AIRIMP. In the '60s, about 60% of BEA's bookings were interline, i.e. through other airlines.

The reservations task did not much lend itself to the prevalent batch processing of sequential files, and was very labour-intensive. Passenger records, for example, were maintained manually in card files of hand-written cards. To complicate matters further, flight schedule changes needed to be managed and, ideally, seat availabilities adjusted, through the months prior to a flight in the light of (not very reliable) statistical booking patterns.

To replicate all this work, with a computer system, gave rise to further requirements. It would need to be 'user-friendly' (a later buzzword), not least for hundreds of telephone sales clerks. It would need high performance – responses to transactions within 2 seconds – much faster than that was considered to be user-disconcerting. It would need to be extremely reliable, and operational 24x7, being crucial to the airline's core business. In fact, the successful Univac proposal and contract specified, by a carefully constructed formula and with attaching financial penalties, better than 99.97% system availability, round the clock.

When the project began late in 1963, there were a few precedents. American Airlines, with SABRE, and Pan American, with PANAMAC, were implementing similar systems with IBM. More relevantly, Eastern Airlines, Capitol and NorthWest had completed numeric or 'inventory'

implementations with Univac, but not full PNR systems. In 1965, having lost at BEA, IBM won a system for BOAC, based on the delusion of a ready-made reservations application package and a serious underestimate of hardware requirements.

The risk and investment cost in these early systems were huge, hardware costs alone amounting to several £millions (in old money). What were the business motives?

There were several, with varying degrees of plausibility, as the 'me too' fashion for such systems spread among the carriers. The original, most measurable motive among the major airlines was cost-reduction, meaning significantly fewer staff in the labour-intensive manual reservations function – which also tended to be the least unionised. This bottom-line benefit was easily measured and predicted.

Better load-factors, i.e. more bums on seats per flight, was another benefit likely from greatly improved inventory control, but less readily quantifiable, because of other contributory factors. So, too, were enhanced customer service and the advertised boast of leading-edge computer technology: expensive, but cheaper than new aircraft.

Hardware

The heart of Beacon's hardware complement comprised two Univac 490 mainframe computers with fixed head and moving head drums amounting to 700 million characters.

The 490 had 32k words of core memory, of 30 bits, addressable as half, third or fifth words, the latter representing a six-bit character when required. The single-address instruction set was relatively limited, but powerful, including several 'replace' commands and the ability to test an arithmetic result in various ways and 'skip' or not the following instruction. Instructions had 15-bit addresses, modifiable by any of seven index (B) registers, while A and Q were arithmetic registers. Add time was 4.8 µsecs, and average execution time was six µsecs. There was no memory protection – seen as an advantage by performance-conscious 490 fans, provided programs behaved properly...

Normally, most of the peripheral equipment was attached to the currently on-line 490 computer. All subsystems were individually switchable between the on and off-line computers by means of manually operated switch buttons.

Over 400 'agent set terminals', of which about 200, in 1967, were installed in nine remote locations, were connected in groups via buffer processors, small minicomputers used as early cluster controllers. The

agent sets, manufactured by Sperry Gyroscope, comprised a teleprinter, keypad and upright keypad/display, onto which an appropriate route schedule slide was inserted; the slide was hole-coded to transmit the route details. The teleprinter component was principally used in the second, passenger records phase of the project. In subsequent reservations systems, and in later Beacon applications, these relatively slow and cumbersome terminals were replaced with more flexible VDUs.

The hardware occupied an air-conditioned computer room of some 6,000 square feet on the third floor of the West London Terminal on Cromwell Road, West London. Most of the local agent sets were located on the fourth floor of the building, in an area of some 18,000 square feet, supported by automatic telephone call distributors.

A round-the-clock team of 15 Univac hardware maintenance engineers watched over the installation, to minimise breakdown times.

Software

The 490 came with REX, Real-time Executive; totally inadequate for Beacon's requirements. The system was therefore built on a home-made Executive, CONTORTS. Other significant project-developed tools included an interpretive on-line transaction trace, a general utilities system (GUS) and an online training mode for terminal operators.

Time

The hardware costs alone were huge; around £3M in 1963 – upwards of £40M in today's money. It was, therefore, important that the benefits were real and measurable, and that the return on investment began as soon as possible. This meant that the initial application, passenger reservations, had a *phased development*. First, the inventory, or numeric control of flights bookings, was developed and cut over to live operation, in April 1965, delivering the benefits of improved load factors and staffing efficiencies. Then the 'PNR', or passenger records processing, together with handling of interline traffic through the SITA network, was developed, and cut over in stages. A third phase extended the real-time service to numerous remote offices in the UK and mainland Europe, and eventually on-line ticketing was added. Subsequently, other related applications were added, on an integrated basis.

This may seem obvious, commonplace in IT projects, even. Back then, it was not. Once the first phase was implemented, it was committed to the core of the airline's business – 24x7. Adding major enhancements securely to such an operational environment is not easy, particularly when, as with PNR, it involves the rolling transfer of millions of existing

paper passenger records, entered by keyboarding, while those records continued to be subject to additions, cancellations and amendments. Then add the requirement to train several hundreds of user staff in the additional features of the system, and to incentivise their contribution of huge amounts of overtime for the keyboarding task

The cutover of the second, PNR phase, occurred in 1966, and was a substantial project in itself, more organisational than technical. It was planned for Easter. Bank holidays were particularly useful, being long, and light in on-line reservations activity; plenty of time for expensive sales clerks' overtime, and sleepless nights for key development project staff.

It failed. The phase 1 implementation was sadly and successfully reinstated for the Tuesday morning. For various reasons, it failed again, every Bank Holiday until the last of the Summer. Even then, subtle but stubborn bugs persisted including one which struck every morning at peak time in the heart of COHORTS which added one randomly to the contents of a word anywhere in core memory.

People

The development project team, of around 25-30 people, was roughly evenly divided between Univac and BEA staff. It was divided into teams responsible for specific parts of the application and system software, from three to six in each. Team leaders came from either company, appointed on the basis of perceived contribution, in terms of skill and effort. Until the reservations development was successfully completed and signed off, Univac retained system responsibility and project management, along with the challenging contractual obligations on performance and availability. The three-phase reservations development took about 80 man-years. This is somewhat misleading, since hours worked would become, for some, quite brutal, or heroic, (according to taste), and not respectful of public holidays, weekends or private life.

Univac's project management was vigorous, some might even say brutal. Mindful of the contractual obligations and of the finite project headcount, it was not tolerant of non-commitment or technical inadequacy, and this sometimes extended to BEA as well as Univac team members. On the other hand, contribution was remembered and proportionally rewarded, sometimes long afterwards, by promotion and/or recruitment (in the case of several BEA staff), and/or recommendation elsewhere.

The philosophy, later articulated (within ICL!), was that the project goal (on-time, to specification and budget) was the only thing which mattered to the project team. Nothing, from fire on the site to strike action to bureaucratic obstruction in either organisation would be allowed to get in its way. If either host organisation found the project's activity disruptive, well, they should not have instigated it; substantial projects are, of their nature, disruptive.

Systems analysts in BEA originally determined the requirements of the reservations system. During the implementation project, there were no specific analysts on the team; everyone was a programmer. (This changed subsequently.) The Univac management, extremely conscious of the trade-off between nice-to-have features and precious computing hardware resource, and of a limited team headcount, did not feel a need for middlemen, brokering the fine-tuning of requirements, between system end users and programmers. Indeed, one manager somewhat insensitively defined 'analyst', before an audience including several of them, as the opposite of 'catalyst': an agent which, while taking no active part in the process, nevertheless impedes its progress.

Design Features

Prepared by the forgoing background information, it remains to attempt to justify the presumptuous title of this memoir; firstly, with some basic design principles; secondly, with a selection of design features – and remembering that this was the mid-1960s.

Design for Failure:

The site's hardware engineering manager once asserted that hardware, of its nature, was bound to fail occasionally, whereas software, eventually, would become bug-free and thus cease to fail. The programmers (software engineers, we would now say) present did not share his confidence. There was too much unpredictable concurrency of program execution in the system to allow the 'provably correct' testing which was starting to be fashionable for batch processing. In any case, the operational and contractual requirements, of greater than 99.9%, 24x7 system availability, required a system design which would recover rapidly and completely from a system 'crash', however caused – within 30 seconds. To this end, the duplication and re-configurability of hardware has already been described – see 'Hardware' above.

This hardware redundancy needed to be complemented by bespoke software features, which, to varying degrees, were programmed into all parts of the application. Embracing them all was SYCOM, an evolving

set of operator commands allowing complete control of configuration and recovery, down to specific groups of agent set terminals or individual SITA network circuits. 'Complete' recovery meant that, for example, all interrupted multi-step terminal transactions would resume from the last successfully completed step, using drum-stored transaction journals. It was many years later that 'standard' operating systems began to offer this degree of resilience, at a time when just re-starting the OS took minutes, anyway.

Usually, crash recovery would involve manual swap-over of online/offline peripheral sets between the two CPUs and restart of each in its new mode. Each would determine its mode and peripheral population by a process known as 'grope'. Normally, the online CPU would have duplicated files on duplicated drum subsystems. If, however, a drum or subsystem were unavailable, the software would adapt. When a missing unit was restored, its file contents would automatically be brought up-to-date – concurrently with the ongoing online operation of the system. This feature also was not 'standard' for 15-20 years, except on a few premium-priced 'non-stop' system products.

For additional resilience, and for fallback in the case of cutover attempts, flight inventory files were backed up hourly onto magnetic tape, and PNR files daily – concurrently with the on-line TP operation, naturally. If a duplicate drum file subsystem was temporarily unavailable, then, following restoration, it would be reloaded from magnetic tape and then brought up to consistency against the surviving copy – concurrently with ongoing TP, naturally.

Design for Performance

A contractual requirement was that over 90% of interactive transaction steps had a response time of less than 2 seconds. It was recognised that most of the existence time in the system of a sub-transaction would be determined by drum accesses, not processing. Such accesses would be, for example, to read and write (normally twice) file records, update transaction journals (for recovery) and retrieve necessary program elements. Moreover, most drum transfers were small, and transfer speeds quite fast.

Most file accesses on Beacon, as in most TP systems, were reads. Normally, read access requests were spread evenly between duplexed subsystems. Writes, of course, had to be made to both.

CONTORTS gave drum subsystem interrupts highest priority, and special efforts were made, even hardware modifications, to reduce

interrupts overall and to minimise interrupt processing under interrupt-lockout. All of CONTORTS, which occupied about 12k words, and all the online reservations application were memory-resident, to reduce accesses for overlays. All this code was 're-entrant', to allow concurrent processing of multiple transactions; the purpose of this was, at peak times, to ensure drum subsystems never had to wait for the next access request. Files on moving head drums were placed so as to minimise head movements.

The passenger files, incidentally, consisted of the indices and the PNRs – passenger name records. They were arranged as what we later learned were 'inverted files', or even a specialised form of relational database. If a PNR involved, say, four flights, then it would be linked to four indices. Indices and related records were linked together in both directions, and each contained minimal data from the other to allow a degree of re-construction in the event of corruption – a feature which was to prove invaluable.

Instrumentation

Every aspect of BEACON's operation which might affect performance or reliability was counted or timed, so that performance was known and could be tuned. For example, not only were drum accesses counted, but also their duration. Transaction existence times were important in determining the optimum level of transaction concurrency.

Interline Message Handling

When Beacon Phase 2 cut over, including handling and generation of interline reservations messages through the SITA telegraphic message network, as it then was, about 60% of BEA's reservations came by this means, i.e. from other airlines, and only two or three airline systems had reached that stage. Thus most messages were keyboarded on teleprinters, by airline or agency staff, following, approximately, the AIRIMP message protocols. The SITA switching centres were then also not computerised, and could not store the (typically 75 bps, CCITT2-coded) messages. So if a reservations system could not receive incoming messages, business was lost.

Beacon had 16 input and four output simplex circuits. Each input circuit terminated in a Siemens receiver-transmitter, which punched and interpreted the messages on paper tape, which was normally then read directly, under system control, into the system. If and while an enabling signal was not received from the system, between each message, for any or all circuits, loops of paper tape would accumulate, preserving the

messages for later processing. Messages were edited, then stored on drum queues for processing.

Input Error Correction

Experience with other early reservations systems showed that conventional editing of interline messages for AIRIMP format conformity produced about 70% rejects, which would then require human intervention; not good for business or productivity. GEDIT was an interpretive editing meta-language, which corrected and prepared queued messages for processing and response. As usual, a compromise was reached between efficiency and hardware resource use. Beacon settled for a rejection rate of about 7%, an achievement greeted with disbelief by visiting executives from other airlines.

Interactive terminal handling

Several hundred Beacon agent set terminals were clustered under the control of buffer processors, small, early minicomputers, which sent and received complete messages to the online mainframe computer. This cluster controller approach, reducing mainframe connectivity, data errors, processing and, above all, interrupts, became a standardised product line about ten years later, e.g. with ICL's 7502, 7503. Even today, many multi-user minicomputers are still directly connecting dumb terminals, on which any keystroke, unbelievably, causes a central processor interrupt.

Name Searches

A Passenger Name Record (PNR) held details of one or more passengers travelling together on one or more flight/dates on one or more airlines, with specific details, e.g. diet or disability, contacts etc. It could be accessed, most efficiently, by a reference number, or else by any name in the party and any flight/date in the itinerary. Hence the 'relational database' structure, as we later learned to call it.

Any call centre, and airline reservations were the first, has a problem with names. Beacon could, if necessary, search a flight/date's passenger list for best phonetic matches (e.g. Knight, Night, Nite, Nought), by character-sequence matching, or, most tediously, list all the flight/date's passenger names alphabetically.

Postscript

And what happened next? Beacon, and Boadicea, added further applications to their systems, until, in the early 1970s, Government decreed that the two national airlines should merge, creating British Airways. This raised the question about future IT strategy. Beacon won

the technical, price-performance arguments, Boadicea won the political debate, in 1973. (That is one version, anyway!) Beacon was phased out. One factor was that Univac abandoned out the 490 Series, which meant that a Univac-based joint system (the 1110 was proposed) would not be program-compatible with either Beacon or Boadicea.

Meanwhile, a global war had developed between IBM and Univac for leadership in the air transport industry, one of the few markets where IBM's dominance in mainframe-based systems was seriously challenged. To a degree, that rivalry lingers still.

In 1987, four leading Univac airline users founded AMADEUS, a business dedicated to providing airline IT services, starting with reservations, to multiple airline users. This concept had its origins in the mid-1960s – too soon. AMADEUS commenced service in 1992, for airlines and also, crucially, for travel agents. It now claims about 150 airline customers, from British Airways and Qantas to Air Congo and Air Montenegro, some 400,000 'on-line terminals', 2,000 'IT experts', network centres near Munich, Sydney and Miami, 99.9% availability and some 12,000 messages per second. It claims market leadership. So, probably, does GALILEO, a rival service founded at about the same time by other European airlines, notably British Airways. Both are seriously challenged by the global expansion of SABRE, which grew out of the pioneering IBM system for American Airlines, and became the most profitable part of the business.

Beacon had a legacy, of questionable value, for ICL, in its efforts to address transaction processing convincingly. In the late 1960s, a group of Beacon veterans formed part of the System D Project, a TP-orientated operating system anticipating the advent of ICL's 2900 mainframe range. It came to nothing. In 1972, following the appointment of ex-Univac Geoff Cross as ICL's Managing Director, a second, larger wave of Beacon veterans arrived including Ed Mack, who became Director of ICL's Product Development Group. VME/B was already established as the multi-purpose operating system for 2900 Series. Mack established a new TP operating system project, VME/K, considerably staffed by people with Univac airline systems experience. For numerous and varied reasons, beyond the author's knowledge or competence to evaluate, VME/K became not a success, but a trauma. VME/K was finally terminated in 1981 by newly-installed Managing Director, Robb Wilmot.

Editor's note: This article is an abridged version of a talk given by the author to the Society at the Science Museum in February 2008. Michael Knight can be contacted at michaelknight242@tiscali.co.uk.

The Ghost of KDF9

Brian Wichmann

On the 29 August 1980, the last two KDF9s were switched off at the National Physical Laboratory (NPL). Unlike some of the first generation of computers, no attempt was made to preserve a complete machine, although some parts were retained as souvenirs.

Around three years ago, it was decided to collect information about the KDF9 computers to ensure that future historians could study the material. This article is by way of a progress report.

We do have much material of great value including:

- Machine code manual
- Binary encoding of the instruction set
- Listing of one of the early Director programs (control program)
- User manual for Algol, etc.

Thanks to the work of Hans Pufal, we have a KDF9 interpreter which can execute part of the Director — this gives the hope that a few modest KDF9 programs could be brought back to life. In fact, it appears we have the listing of three versions of the Director.

On the negative side, even the list of machines sold is incomplete, but currently is in Table 1 below.

We also lack logic diagrams for the hardware and only few programs have been saved — but this includes the Whetstone compiler/interpreter and part of the Kildgrove compiler. Unfortunately, Leeds University did not retain a copy of the Eldon2 operating system, as was planned.

We have one resource in abundance — memories of many people who worked on various aspects of the KDF9. The heroic story of the two engineers who eventually got the Birmingham machine to work in June 1964 is something I will remember since I always thought the first KDF9 was working in 1963! I started work at NPL in October 1964 which implies I was not far behind. . .

We would very much like to hear from anybody else of their experiences with KDF9 (contact: Brian.Wichmann@bcs.org). The initial group of those contributing to the collection of information has been: Bob Beard, Roger Broughton, Bill Findlay, Alan Freke, David Holdsworth, David Leigh, Roderick McLeod, Michael Poole, Hans Pufal, Peter Stanley, Graham Toal and Chris Willis.

For those who want to learn a little about the machine, the entry in Wikipedia is to be recommended (some other Internet sources have errors). We have yet to decide how to present the information we now have to hand; however, Hans Pufal has already collected a lot of useful information at www.pufal.net/KDF9.

<i>Machine</i>	<i>Working From</i>	<i>Working To</i>
Admiralty Research Lab	?	
Baric (service m/c)	?	
Birmingham University	21 Jun 1964	1972
Bristol Siddeley Engines (later Rolls Royce Bristol)	1964	> 1977
Bristol Siddeley Engines (ex Bristol Aeroplane Co.)	1969	> 1977
Bristol Aeroplane Co.	1964?	1968
Culham – Atomic Energy Authority	Mar 1965	1970/71
Edinburgh University(temporary)	1968	
Glasgow University	1964	1972
ICI 1 (Teeside)	Feb 1964	?
ICI 2	?	
Knutsford Nuclear Power Group	?	
Leeds University	1964	1976
Liverpool University	1964?	
Met office	< Jul 1965	
Newcastle University	1964?	Aug 1973
Nottingham University	1965	End 1971 – mid 1972
National Computing Centre (NCC)	?	1970
NPL 1	01 Aug 1964	29 Aug 1980
NPL 2 (ex NCC)	Feb 1970	29 Aug 1980
Oxford University	Jan 1965	Dec 1971
Salford University	1964?	
Sun Life	?	
Sydney University	1964	
Wills Tobacco	1965	1968?
Winfrith – Atomic Energy Authority	Dec 1964	

Table 1 Details of KDF9 machines traced as of June 2008

We have not yet started to work on the long-term presentation of the material we have collected, but this clearly needs to be integrated with the other CCS material and also with material about KDF9 from other sources.

Letters to the Editor

From Hugh McGregor Ross

Dear Nicholas

Just a note to add to the many exclamations of praise you must be receiving for your editorship of *Resurrection*. Our paths have crossed several times in that connection, always with great harmony. In particular, you provided the only way in which records of the pioneering British computing work could be recorded, and made available to a wider audience.

I most especially appreciate your final issue (42), just received, for a particular reason. It combines, and shows up the difference between, two entirely different and even contrasting ways of doing that recording.

The article by Peter Barnes, while touching on many different topics (yet without making any mistakes) is redolent with comments that convey the extraordinary, and now inevitably lost, atmosphere of working at the start of the computing revolution. I was the man in Ferranti encouraging de Havilland, aiding Peter in their purchase of a Pegasus computer and providing some of the support to make it useful.

Simon Lavington's article, on the other hand, is intensively researched, is of impeccable factual accuracy, and provides much information that is not otherwise available. Never before has such a complete review of an early computer company's achievements been presented.

The difference is well illustrated by Simon's statement:

"The NRDC/Elliott 401 computer was exhibited at the Physical Society Exhibition in London in 1953, when it caused a stir."

Certainly it did that, but what about the background?

I was well aware of the significance, at that time, of exhibiting new technological products at that annual exhibition. Several new items that I had developed while with my previous employers, Kodak Research Laboratories, had been exhibited there in earlier years. These demonstrated, at a very early stage, the application of negative feedback to electro-mechanical equipments to enhance their performance and stability.

Bill Elliott was determined to gain recognition for his concept that stored-program digital computers could be designed around a small

number of different types of logical unit, and could be made to work satisfactorily by implementing those logical units as plug-in packages that could be manufactured and tested independently of the complete computer. It was a revolutionary concept, and attracted much criticism.

To achieve this goal by exhibiting the 401, he had to move heaven and earth to get the machine finished in time, and he inspired his team of colleagues, which may well have numbered no more than ten, to work with a kind of obsession, as he did.

In the event, he demonstrated: that such a computer could be designed and made; that it could be dismantled sufficiently to be transported; that it could be re-erected and made to work in the short time available; and that it could do useful work.

That is why it caused a stir, and it was well justified.

It was only a few years later that I followed his example. By that time the work that Bill Elliott initiated while at Ferranti had resulted in the initial version of what we called Pegasus II. This took the form of providing an effectively working magnetic tape system, together with the necessary software, that was sufficiently error-free to permit its use for storing, and updating, the enormous data files that lie at the heart of commercial data processing applications.

I was responsible for putting a Pegasus II into an exhibition for automated data processing, and I still have a photograph of it, with two men engrossed in controlling the machine at its console, and it is possible to see that one of the four magnetic tape units is actually working!

Very sincerely Hugh

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.

Obituary: David Caminer

Frank Land

David had joined the tiny LEO group to establish and develop a team which would be able to realize Lyons' ambition to utilise an electronic digital computer to streamline the already efficient office work of the organisation. This meant designing robust systems and programming these onto the computer. Nothing like this had been done anywhere in the world before. David's achievement was to carry this through successfully. His special skill was in understanding both the business tasks the computer was to be employed on and the capabilities – or their lack – of the primitive computer we then had available.

Much of the success of the early applications stem from the high standards David set for himself and his team. The standards applied to every element of the task of bringing an application from concept to reality. He did not spare himself, and he expected the same from those who worked with him. If a document was handed to him by one of us which was poorly phrased or not clear, it was sent back until it satisfied the highest standard. He could be a hard taskmaster. But those of us who survived – and most of us did – became some of the standard bearers world wide for the new profession of Systems Analysts.

Despite the many changes in the British computer industry as companies merged, and LEO became absorbed by a succession of new industrial corporations, David never gave up. He gave his loyalty and skills to each successor company until his retirement. He earned his OBE for bringing to a successful conclusion a major project for the European Union.

But if all his endeavours were devoted to serving his employers, his love was reserved for LEO. This is exemplified by the effort he made in his retirement to chart the history of LEO. The progress of LEO through three decades is documented in the book he edited and inspired. It is now one of the classics of computer history. Later he founded the LEO Foundation as a vehicle for keeping the spirit of LEO alive.

David never lost that spark or his striving for achieving the highest possible standards. David was one of the heroes of the computer age. His was a life to celebrate.

A longer and more detailed appreciation of David will be published in a later issue of *Resurrection*.

Obituary: Professor Hilary Kahn

Dik Leatherdale

Professor Hilary Kahn of Manchester University's Department of Computer Science died in November after a long illness.

Hilary joined the Department in 1967 and applied the then fairly new discipline of CAD to the development of the MU5 computer system. This experience was then used to found the Department's CAD research group, a connection she maintained until retirement in 2006. Springing from her interest in CAD, she also applied herself to research in information modelling.

However, she will be best remembered by CCS members as the main organiser of Computer 50, Manchester's 1998 celebrations of the 50th anniversary of the world's first stored-program computer (helped by husband, Dr Brian Napper, who created the website www.computer50.org). Those of us who had the privilege to be there (we filled the Bridgewater Hall, no less!) had little idea at the time of the massive amount of time and effort that Hilary had spent on the event. It was, in every sense, a triumph.

Your editor can personally testify to Hilary's kindness and generosity. She is much missed.

Forthcoming Events

Every Tuesday at 12:00 and 14:00 Demonstrations of the replica Small-Scale Experimental Machine at Manchester Museum of Science and Industry.

Every day Guided tours and exhibitions at Bletchley Park, price £10.00, or £8.00 for children and concessions. Exhibition of wartime code-breaking equipment and procedures, including the replica Bombe and replica Colossus, plus 60 minute tours of the wartime buildings. Go to www.bletchleypark.org.uk to check details of times and special events.

11 September 2008 London seminar "The Legacy of LEO in Business Applications – a Tribute to David Caminer". Speakers Frank Land and the LEO Society.

- 16 October 2008** London seminar “A Life with Databases”. Speaker Charles Bachman.
- 21 October 2008** NWG seminar “Computing Before Computers”. Speaker Jane Wess.
- 13 November 2008** London seminar “Bull Computers”.
- 18 November 2008** NWG seminar “Elliott’s Contribution to Computers”. Speaker Laurence Clarke.
- 15 January 2009** London seminar “ICL-Fujitsu Technology Collaboration”. Speakers Tom Hinchcliffe & John Vernon.
- 20 January 2009** NWG seminar “Some CAFS Applications”. Speakers Hamish Carmichael and Martin Wright.
- 17 February 2009** NWG seminar “Computers in Telephony”. Speakers David Parsons and Nigel Linge.
- 19 February 2009** London seminar on “EMI Computers – A Reminder”. Speakers Michael Knight & John Bradbury.
- 12 March 2009** London seminar “The Archaeology of Very Early Algorithms 1948-58 : Strachey’s Love Letter Generator”. Speaker David Link.
- 23 April 2009** London seminar “JANET – the First 25 Years”.
- 14 May 2009** London seminar “BBC Domesday Book Project”.

Details are subject to change. Members wishing to attend any meeting are advised to check the events page on the Society Web site at www.computerconservationsociety.org for final details which will be published in advance of each event. Details will also be published on the BCS Web site (in the BCS events calendar) and in the Events Diary columns of *Computing* and *Computer Weekly*. London meetings take place in the Director’s Suite of the Science Museum, starting at 14:30. North West Group meetings take place in the Conference Room at the Manchester Museum of Science and Industry, usually starting at 17:30; tea is served from 17:00.

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. Queries about Manchester meetings should go to William Gunn at william.gunn@ntlworld.com.

Committee of the Society

Chairman: **Dr David Hartley FBCS CEng** Email: david.hartley@clare.cam.ac.uk

Vice-Chairman: **Tony Sale Hon FBCS** Email: tsale@qufaro.demon.co.uk

Secretary, Chairman DEC Working Party: **Kevin Murrell** Email: kevin@ps8.co.uk

Treasurer: **Dan Hayton** Email: Daniel@newcomen.demon.co.uk

Science Museum representative: **Dr Tilly Blyth** Email: tilly.blyth@nmsi.ac.uk

TNA representative: **David Glover** Email: david.glover@nationalarchives.gov.uk

Bletchley Park volunteers representative:

Pete Chilvers Email: pete@pchilvers.plus.com

Chairman, Elliott 803 Working Party:

John Sinclair Email: john@eurocom-solutions.co.uk

Chairman, Elliott 401 Working Party:

Arthur Rowles Email: rowles01@globalnet.co.uk

Chairman, Pegasus Working Party:

Len Hewitt MBCS Email: leonard.hewitt@ntlworld.com

Chairman, Bombe Rebuild Project:

John Harper Hon FBCS CEng MIEE Email: bombe@jharper.demon.co.uk

Chairman, Software Conservation Working Party:

Dr Dave Holdsworth CEng Hon FBCS Email: ecldh@leeds.ac.uk

Digital Archivist & Chairman, Our Computer Heritage Working Party:

Professor Simon Lavington FBCS FIEE CEng Email: lavis@essex.ac.uk

Editor, Resurrection: **Dik Leatherdale** Email: dik@leatherdale.net

Web Site Editor: **Alan Thomson** Email: alan.thomson@bcs.org

Archivist: **Hamish Carmichael FBCS** Email: hamishc@globalnet.co.uk

Meetings Secretary: **Dr Roger Johnson FBCS** Email: r.johnson@bcs.org.uk

Chairman, North West Group: **Tom Hinchliffe** Email: tom.h@dsl.pipex.com.

Dr David Anderson Email: cdpa@btinternet.com

Peter Barnes FBCS Email: panda.barnes@btinternet.com

Chris Burton CEng FIEE FBCS Email: cpb@envex.demon.co.uk

Professor Martin Campbell-Kelly Email: M.Campbell-Kelly@warwick.ac.uk

George Davis CEng Hon FBCS Email: georgedavis@bcs.org.uk

Peter Holland Email: peterholland@care4free.net

Dr Doron Swade CEng FBCS Email: 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 37 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