

Computer Conservation Society

Aims and objectives

The Computer Conservation Society (CCS) is a co-operative venture between the British Computer Society and the Science Museum of London.

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

The aims of the CCS are to

- ◇ Promote the conservation of historic computers
- ◇ Develop awareness of the importance of historic computers
- ◇ Encourage research on historic computers

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

The CCS is funded and supported by a grant from the BCS, fees from corporate membership, donations, and by the free use of Science Museum facilities. Membership is free but 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.

The corporate members who are supporting the Society are Bull HN Information Systems, Digital Equipment, ICL, Unisys and Vaughan Systems.

Resurrection

The Bulletin of the Computer Conservation Society

ISSN 0958 - 7403

Number 8

Winter 1993

Contents

Editorial

Nicholas Enticknap, Editor 2

Guest Editorial

Peter Hall, Chairman, North West Branch 3

The CCS branches out

Dan Hayton, Treasurer 4

Obituary: Ted Newman 6

Obituary: Bernard Swann

Hugh McGregor Ross 7

Early computer development at NPL

Donald Davies 8

The coming of mix and match hardware

Robin Shirley 16

NRDC's role in the early British computer industry

John Crawley 25

Working Party Reports 33

Letters to the Editor 36

Forthcoming Events 40

Editorial

Nicholas Enticknap, Editor

There have been two major developments since the last issue. First, the focus of the Society's activities has now switched to Blythe House. The Elliott 803 and the various DEC systems made the short journey from the Old Canteen building at the Science Museum in early November, and the working parties are currently reassembling their systems and preparing to start restoration work.

The Elliott 401 stays in the Old Canteen building, and will join the 803 in the Elliott room at Blythe House at a later stage. The Pegasus also remains *in situ*—we still await a decision on its future location.

The second development is that the new North West Branch is up and running. An inaugural meeting held in late October was an unqualified success, with 60 people coming along to listen to Peter Hall, Frank Sumner and Charlie Portman give their accounts of the “Challenge of the Fifties”. Dan Hayton reports on the meeting on pages 4-5, while Branch Chairman Peter Hall is the contributor of this issue's Guest Editorial.

Sadly, Peter's first duty at the meeting was to give the news that Liz Segal, who played a major role in setting up the new branch, had been taken seriously ill and has had as a result to resign her position as Secretary. All at the CCS send her our good wishes for a speedy recovery.

Better news is that, after two years of prolonged negotiations, it now appears that the battle to secure Bletchley Park as the home of a proposed Museum of Computing and Cryptography is entering its final phase. The Coopers & Lybrand feasibility study on the future of the World War Two code-breaking centre, which endorsed the Bletchley Park Trust's original proposals to Government, has now been accepted by all relevant parties. We expect a speedy resolution of the remaining details under negotiation.

The feature articles in this issue are a mix of ancient and modern. John Crawley writes about the role played by the NRDC in the development of the early British computer industry, while contemporary events at the National Physical Laboratory are chronicled by Donald Davies. The modern touch is provided by Robin Shirley, who writes about the influence of the S-100 bus on the development of the microcomputer industry in the seventies.

Guest Editorial

Peter Hall, Chairman, North West Branch

The North West lays claim to the birth of the Industrial Revolution, and we in the North West believe we have a good case to be at least one of the birthplaces of the computer. Much of the pioneering work took place in this part of the world, and many of the people involved still live in these parts. The Manchester University team, of course, springs immediately to mind, but there was also much pioneering work at Ferranti (Sirius and Orion for example), English Electric (eg KDF9 and System 4), and AEI (eg 950 and 1010). I hope that any members of the Society who are able to will join our ‘provincial’ activities, particularly if they can bring some ‘know-how’ on ‘our’ machines! (Please contact me on 0260 224363.)

By the time this note is in print we will have had our first meeting. I hope that we will have been able, at that meeting, to excite interest and a realisation of the importance of our work by talking about the “Challenge of the Fifties”. Who knows, by the time you read this, we may be overwhelmed with ideas and volunteer workers! Reading the back issues of *Resurrection* I became conscious of the very real problems faced by the Society. There is so much that could be done but resources, both human and financial, are limited. How do we settle priorities?

To add to the problem I would like to suggest one area which has not been significantly addressed as yet. How about those long suffering people, our pioneering users? Should we not do some work on the early users of computers in banks, insurance companies and so on? The struggles to make those early unreliable machines pay off are surely worth recording.

Looking back from where we are today and the struggles we are having to conserve and archive the early days, surely indicates that we should also be setting up mechanisms now to determine what of today’s artefacts (hard and soft) are seminal and need therefore special treatment now rather than in 10 years time.

Finally I have to admit the reason for agreeing to get involved. Two Christmases ago my nine year old grandson was given a computer. He demonstrated what looked like considerable expertise on this machine (well in excess of my own), so I asked him if he had a computer at school. “Yes”, he said, “256K — useless!” What an illustration of the rate of change, and of the need urgently to record those days when 16K was a luxury!

The CCS branches out

Dan Hayton, Treasurer

Thursday 21 October saw the inaugural meeting of the North West Branch, held in the Goldstone Room of the Manchester Museum of Science and Industry. The Museum's interest in the CCS, and ours in their activities, had been awakened by a number of preliminary visits by Tony Sale, Chris Burton and myself to examine their Pegasus and determine what could be done in the way of conservation. It also seemed that, given the links between Manchester University, local companies and the earliest origins of stored program electronic computers, Manchester and the neighbouring area offered great potential for research into the people as well as the machines themselves.

As an "outsider", I was delighted to see many acquaintances being renewed over coffee, and was pleased to be asked to fetch extra seats. In addition to existing members of the Society, many newcomers, ranging from pioneers of computing to students, soon used up the supply of membership forms available.

Peter Hall, chairman of the North West Branch, opened the meeting with the sad news that Liz Segal, who had taken on the role of secretary of the new branch, had been in hospital and had, with regret, tendered her resignation due to ill health. The meeting expressed its thanks to Liz for her efforts in setting up the branch, and its good wishes for a speedy recovery.

Peter then introduced Graham Morris, Tony Sale and myself from "down South" and Jenny Wetton of the Museum. Following a welcome by Graham on behalf of the Committee, Jenny outlined the plans for the Museum's activities around the conservation and display of its Pegasus, and in collecting the oral history of the computer industry.

Peter Hall then introduced Frank Sumner, who gave us a view of programming the Mark I both from the theoretical and practical points of view.

A reprint of the second edition of the programming manual provided an illustration of what could be described as the original risc system—some 16 instructions were all it had—and training, which consisted of being given the manual to read. The practical difficulties of running early electronics next to a tram line were compounded by printers which could at any moment throw pieces of their mechanism right across the room—this

explains why the engineers kept an army surplus tin hat handy!

Frank presented a copy of this edition of the manual to the Museum. He said he did not have access to a copy of the original edition, and appealed to anyone who knew of the existence of one to get in touch.

Charlie Portman then spoke of the design and development of the series of Ferranti machines, covering the design, manufacture and installation, and including the development of software to keep track of the components and wiring of the increasingly complex hardware.

Peter Hall recalled his days managing this new money-eating enterprise, and of being introduced to members of the Ferranti family as one of the men who was spending their inheritance.

All the speakers received an enthusiastic response from the floor. Questions ranged from the financial details to the development of software simulators for the machines which did not survive. Chris Burton does not live too far away so there is “local” expertise on hand. Other audience response included a plea for the collection of current equipment and an offer of storage space from a local microcomputer shop.

The meeting adjourned to a pizza restaurant (a long-standing CCS tradition). Those of us from the South—Graham Morris, Tony Sale, Len Hewitt and myself—had unfortunately to forgo this pleasure, making our ways home by train and microwaved burgers.

The most encouraging start has already borne fruit, and Jenny has reported a number of volunteers for the conservation and oral history projects. We all wish the North West Branch success, and look forward to drawing on the speakers it finds as well as sending some of our performers, in person or recorded, to entertain them.

Your chance to help

The creation of the North West Branch gives people who live in the area and are interested in the history of computing a unique chance to become involved. Volunteers are particularly wanted for the Pegasus conservation work described in Dan Hayton’s article, but help of all kinds is welcome, and no specific expertise is necessary. Anyone who would like to play a part in the branch’s activities, or who would like simply to become a member of the Society, should contact Branch Chairman Peter Hall on 0260 224363.

Obituary: Ted Newman

We regret to report that Ted Newman died on August 7. We have lost a good friend, for Ted was an enthusiastic and committed member of the Society from the outset.

He gave two presentations to the Society. The first, in February 1991, described the influence of AD Blumlein on electronic circuit design. The other came as recently as last May as part of the all day NPL seminar, describing the engineering history of the Pilot ACE.

Ted's contribution to the Society, though, extended far beyond his willingness to talk about his own involvement in computing. He was always available to help in arranging seminars and assisting other speakers, and was generous with his time in providing the behind-the-scenes expert knowledge without which a Society like ours could not flourish.

Ted Newman was himself a distinguished contributor to the development of computing in this country. He first made his mark at EMI Research Laboratories, which he joined in 1941 at the age of 22 as "Blumlein's personal dogsbody", as he himself modestly put it in that February 1991 talk. He worked with the legendary electronics pioneer on the design of electronic circuitry for use in military radar systems, including the celebrated H2S airborne system.

After the war this experience was put to use in the design of circuits for television equipment. Newman's idealism led him in 1947 to leave industry for Government service, and he joined the National Physical Laboratory to work on the ACE computer project. His experience of radar and television circuitry proved invaluable in conquering the many computer engineering problems that arose with this pioneering machine. In particular, he is credited with getting the memory to work (in collaboration with David Clayden) and with producing a much improved logical design for the central control unit.

Newman later became involved in work aimed at developing computers for data processing, rather than scientific, applications. With Michael Wright he produced a seminal report for the Treasury which stimulated the use of computers in Whitehall.

As the use of computing and the computer industry developed, Newman remained an influential figure. During the late sixties and early seventies, he chaired the NPL committee which vetted new proposals for Government funding under the Advanced Computer Technology Project.

Obituary: Bernard Swann

Hugh McGregor Ross

Bernard Swann was in my view responsible, more than any other individual, for initiating the transition from computers as an academic, scientific and military activity to being a business and an industry.

Bernard Burrows Swann first achieved distinction as a statistician, particularly in Government and during the Second World War with the Army in India. His ability eventually elevated him to the position of Assistant Secretary of the Statistical Division of the Board of Trade.

Swann recently told how one day, walking along Whitehall, he met Vivian Bowden who had been asked by Sir Vincent de Ferranti to make a commercial success of an engineered version of the first Manchester University digital computer. Bowden invited Swann to join him.

It must have taken extraordinary courage to leave a senior position in the Civil Service and enter computing, which at the time — 1952 — was not a business or an industry, had no obvious prospects, was not recognised as a profession and was of no repute.

Swann's earlier experience with punched card automation convinced him that the future of computers lay in teaching people in industry and commerce to use them. So the Ferranti Computer Centre in London, the first of its type, was set up and equipped with the original Pegasus computer (the actual machine that the Society's new North West group is now restoring).

Here simplified programming methods, subroutine libraries, extensive documentation and manuals, training courses and hands-on experience of using both Pegasus and the Manchester machine were rapidly expanded. This was way in advance of anything else either in industry or in the universities, and was the precursor of everything we now know of as computing activity.

In building up the marketing and software development activity for Ferranti computers, Swann's style of management resulted in a harmonious and intensely dedicated team, many of whom still feel a personal respect for him.

In recent years Swann took a great interest in the recording of early British computing activity. This has had a significant influence on the work being undertaken by the Society at the Science Museum.

Early computer development at NPL

Donald Davies

This article describes the influence that Turing had on the history of computing, and in particular on the work done at NPL. It then discusses the early stages of the NPL computer developments.

Turing and computing theory

Alan Turing is recognised as a genius who made fundamental contributions to the foundations of mathematics. Later he also made very important contributions to code-breaking in World War 2, recorded in a magnificent biography by Andrew Hodges, and in a play, but here I am concerned with mathematics and Turing's significance in that area.

Turing was concerned with the Entscheidungsproblem, which just means the decision problem. It was called that by David Hilbert, who proposed it at an international conference in 1928. Hilbert had already done a lot of work on the foundations of mathematics, but he realised that there were three things that still remained to be done.

They were to address three questions relating to any particular brand of mathematics described in formal terms — a requirement already well understood following Russell and White's publication of "Principia Mathematica". They were the questions of completeness, consistency, and decidability.

The first question, completeness, required that any well formed assertion in a formal language could be proved either true or false. That would be very nice if it could happen.

The second question concerned consistency. If you can prove both a given theorem and its opposite then you can deduce everything, so a system of mathematics that isn't consistent is pretty useless. Whereas completeness is desirable, consistency is essential, and it would be nice to be able to prove it for a particular brand of mathematics.

The third, decidability, required that, given an assertion in some formal language that you can put into a machine or into some formal process that has been predefined, and which is finite in extent, you know it will eventually reach a decision, either that the assertion is true or that it is false.

Those were the three questions that Hilbert posed. Unfortunately for him Kurt Godel at the same 1928 conference produced his incompleteness theorem which destroyed the first requirement. It showed that if a brand of mathematics was both consistent and also contained enough to be able to generate actual numbers, then it could not be complete: there must necessarily be assertions that you can write down which can neither be proved true nor proved false.

This left the question of decidability, which had become a slightly different question because assertions were now in three categories: provably true, provably false, and unprovable.

I don't claim to know exactly how Turing's work related to the Entscheidungsproblem; it certainly made a big dent in it. Turing was concerned to show that there were numbers which could be defined perfectly well but never be computed.

To him a formal machine which would in a finite number of operations arrive at a decision had to be explained; you had to say what the formal machine was before you could begin to prove theorems about it.

So with regard to decidability Turing had to come up with a way of describing formally what any conceivable machine or procedure would look like. In 1935 he had some ideas, which he developed in a paper submitted to the London Mathematical Society and published in 1937, called "On computable numbers with an application to the Entscheidungsproblem".

The title tells what his real interest was, which was to decide just what was a machine was capable of doing. To do this he first had to define a machine. This became the Turing machine, which is now used as the model for a computing process in all work on the foundations of mathematics.

The Turing machine looked a bit different from a present day computer. For one thing it could not use an address to access its store because that would mean the memory had to have a finite size, and although the machine was finite you don't know how big it's going to be. So to make a general purpose Turing machine you had to have access to an unlimited amount of memory — not infinite but unlimited.

So he used a tape with one end cut off and the other end going as far as you need for the process. There were symbols on the tape, and a finite state machine which could read a symbol, look up a configuration table which told it for that symbol and the state it was in what the next symbol and the next state would be, or indeed whether to move the machine left or right for the next operation. That's all it had to do. It was a very

simple machine to describe and it became known as the Turing machine. Let's call it T.

Now that's not enough for our purpose because we have to be able to tell the machine what to do. This was the important step that Turing made.

So far I've talked about configuration tables in the abstract. Turing designed a second machine with a specific configuration table which he called U, the universal Turing machine or universal computing machine.

It is such that when you take the configuration table of T, write it out in the language he described, put it on the tape in symbols which he defined and then start the machine in the right way, it would chug about and (in a very long time) perform an emulation of what machine T would have done. This is really a programmed computer, using interpretive code. In order to define this machine with its enormous number of states, Turing introduced abbreviated tables which were in effect functions or "macros" which would have to be expanded when you came to build the machine.

One of the consequences is that the numbers that you can compute with this machine must be enumerable. They must be in a one to one correspondence with natural numbers, because the result is entirely defined by what appears on the tape, and that is just a number.

The real numbers are non-enumerable but it still might be possible to compute all the ones we can define, because they might be enumerable. Turing's paper shows that there are numbers that can be defined but which cannot be calculated by machine U, no matter what 'program' is written on the tape.

The paper not only made a big dent in the Entscheidungsproblem but it also produced a definition of a computer, and of what it means to have a process which can actually be computed.

The machines T and U were not meant as practical computers but were the means of giving a concrete form to the 'finite, formal process' that is at the heart of Hilbert's concept of decidability. This work was part of Turing's lifelong interest in what a computer could and could not do. This led to his involvement in computers and in particular to his work at NPL.

The origins of the ACE Pilot Model

During World War 2 there were big computational projects in a number of places in the UK. You can't do complex engineering design without

computers as we now know, but the only computers we had were the two-legged type, who had the aid of Marchants and Brunswega calculators.

Probably the biggest team was at the Admiralty Computing Section at Bath. There was also one working on armaments at Fort Halstead. I was involved with one working for the Tube Alloys project on nuclear weapons. We had about a dozen people in one particular team just doing a single computation solidly—I think it was the critical size problem.

It was very much a hands-on type of numerical computation, and great fun, but the need to mechanise it was clear to everybody. Towards the end of the war a number of meetings took place with the idea of setting up a centre for this kind of work after the war.

This centre was discussed at some length and eventually it was decided to put it at NPL, where it became the Mathematics Division. Its primary task was to understand and develop numerical methods and then hopefully with computers (only a distant prospect) to carry these out more effectively.

The initial work at NPL was done using punched card machines and NCR accounting machines. You had to bang on these massive keys like a medieval bell ringer in order to make the machine work. One of the chores for people when they first joined Maths Division was to spend an hour a day, whoever they were, banging the National Cash machines.

Mathematics Division was formed in 1945. JR Womersley became Superintendent. I think he was a much misunderstood man. He was no mathematician: I think it was Wilkinson and Fox who had a bet on whether they would ever come out of his office having seen Womersley write an equation. But he was an extraordinarily good manager, adept at fighting for his cause. Mathematics Division got a reputation for winning battles with Administration by the method which works in the Civil Service, which is to use the rules to win your case.

It was Womersley, by the way, who coined the name ACE, meaning Automatic Computing Engine, in deference to Babbage who used the term engine.

Womersley started to recruit people. One of the first was Alan Turing who joined in October 1945, the Division having been formed in April. He was recruited as a Senior Scientific Officer at £800 per annum.

Turing began work immediately and by December (a remarkably short time) he had produced a magnificent report on the design of ACE. (The NPL has republished this report.) It's quite extraordinary; for example

at the end there is a superb description of the Williams tube, with all the practical and engineering problems described, ending up with the comment: “none of this involves any fundamental difficulty but no doubt it will take time to develop”.

Surprisingly this report never discusses which type of memory to use (it does say at one point mercury delay lines are the system we will probably use, but he doesn't say why he decided that). In his table of characteristics of stores the Williams tube comes out top. (By the way, he *is* talking about the Williams tube, not the iconoscope which von Neumann discussed. RCA did a lot of work on the iconoscope as a memory device but the Williams tube was the winner and Turing was somehow aware of this.)

The report contains all kinds of other details including the physics of electrical delay lines and a lot on valves. What he writes on valves is less impressive because he seems to miss where the main source of the delay in valve circuitry will arise.

This was the most rapid part of the whole development of the ACE Pilot Model. After that it went rather slowly.

Womersley, in his inimitable way, sold the project to the Executive Committee. The decision was delayed twice asking for more information, but by May 1946 (not a bad delay) it had been agreed. Womersley proposed to have a Pilot Model stage which Turing was quite in favour of.

At this time, Charles Darwin, grandson of the great Darwin, was Director of the NPL. Darwin placed a three year contract with the Post Office at Dollis Hill in June 1946. It is likely that the decision to have the development of the ACE done at the Post Office research station was due to their great success in building the Colossus machines for code-breaking at Bletchley Park.

Dollis Hill offered to give it high priority, and they put Coombes and Chandler (formerly of the Bletchley Park Colossus team) onto this project. At the same time Jim Wilkinson, the first of the employees other than Turing to be in on the Ace project, joined NPL. Mike Woodger came later, and next myself and Gerald Alway.

Wilkinson found when he arrived that the design had already reached version 5. The initial 2-address machine (which is what the Pilot Model and Deuce eventually became) had by that time become a 3-address machine. In this the instruction typically said: take something from address A, something from address B, do some function on them and put them in address C. But the emphasis was still on having simple hardware and

rather difficult programming to get the high speed. That was not the philosophy, for example, at Cambridge.

At that point something curious happened. Both Darwin and Womersley were involved in writing to Manchester and Cambridge where other projects had started, asking them if they would design and develop ACE. This is very difficult to understand. Nothing came of these approaches because they were too busy with their own projects.

Another interesting development (which had later repercussions) was the arrival of Harry Huskey. He was a very colourful individual and had worked on ENIAC, so he was a great man to have with us in this project. He had left Moore School and was going to join National Bureau of Standards but he decided he would have a sabbatical year at NPL.

He seemed to be the most lively person in the whole group, really a go-getter. When he first joined he was given the task by Turing of going round to the other projects and the Post Office and making a report. This report still exists.

Huskey wanted to build a Test Assembly himself, and after a lot of pushing he believed he had got permission to do so. I'm not sure that he did. I think that Womersley connived with him in getting something started, but I don't believe it was an official project.

When I joined I found that the version had reached 7C. It stopped there because Turing left at that point. He left a very detailed and careful design for the full scale ACE. I soon became very frustrated with working in the maths team: what we were doing was writing programs for a non-existent machine which we could never test. So after a while I joined Huskey's team and started to work on the Test Assembly.

At about the same time, Ted Newman and David Clayden joined from EMI. They formed our electronics expertise, of which we were badly in need, bringing with them the principles of electronic design from Blumlein who was a genius in that field. They adapted their experience in television and radar to computing circuits and devised a circuit technique. Even today semiconductor analogues of this principle are being used for the fastest logic—it was common cathode logic, nowadays common emitter logic.

While working on Huskey's design, I got a visit from Ted who told me in forthright terms that we were wasting our time using the wrong kind of logic, and we should do it a different way. I agreed with him: I could see that what he was proposing was far better. But it wasn't very long before

Huskey left, and his project was dropped.

At about the same time, an Electronics Section was formed and the Post Office contract was closed. This indicates, I think, an understanding by Darwin that having a remote unit doing the design for us just wasn't going to work.

The electronics section was formed within the radio division of NPL. HA Thomas from the radio division was put in charge. That group was building the infrastructure for the computer, getting the power supply arrangements and working out methods of generating logic circuits, but not actually building equipment that would be the final Pilot Model, except for the delay lines.

The Post Office continued to work on a computer based on the ACE design and we went over from time to time to help them. I remember producing a design for a multiplier, not the best multiplier but it worked moderately well. Later on I was amazed to see that multiplier in the flesh at Malvern, because this machine became Mosaic, which went to Malvern and worked for the Ministry of Supply there.

Our problems in those early days were ones of organisation though we had plenty of good people. We had one team working *in vacuo* on the logical design and another team doing electronics. Things finally came together in mid-1948 when we moved over from the maths enclave in the south of the laboratory site and joined the electronics team in Bushy House. I remember Jim Wilkinson actually cycled over with an oscilloscope balanced on the handlebars of his bike.

Immediately we mathematicians started practical work. I organised an assembly line of mathematicians, after teaching them to solder, to produce a small model of the part of the machine that I thought was going to create a lot of trouble. This was the input staticiser and output dynamiciser. That combined team worked extremely well and produced the machine by November 1950. By then it could be demonstrated to the press and somewhat haltingly do jobs. Some time in 1951 it became a really useful machine.

I want to say something about some of the outcome of Huskey's contribution. When he left NPL he went back to Washington, but he didn't stay there very long: after a conference at the Institute of Numerical Analysis at UCLA in Los Angeles (associated with NBS), he decided he wanted to live out there. So he promoted the idea of building a computer at UCLA.

This was SWAC, Standards Western Automatic Computer. It was

nothing like the Pilot ACE, but some time later Bendix Corp approached him and said, “Look, we’d like to make a small and economic computer that would sell in large numbers”. Huskey based this computer entirely on the Pilot Model design, using magnetic drums.

He made them into recirculating stores by having them read the information at one point and then rewrite it at another point. In that way he could wrap a number of delay lines round one track of the drum and make short delay lines if he wanted to. So he had an exact analogue, although eight or nine times slower, of the ACE, working with magnetic drums. This was called the G15 computer, and 400 of them were built, so in terms of the outcome in hardware it was much the most successful of the ACE analogues.

The last G15 had a gold plated control panel. As a memento Bendix gave Huskey one of these machines which he installed in his garage. I went to see him once: he opened the garage door and a mass of heat came out. It was warm outside but even warmer inside. I believe Huskey had the first home computer. That machine is now in the Smithsonian Museum.

There was another outcome I was reminded about by RT Clayden, David’s brother, who left NPL, went to EMI and built the EMI business computer for them. This was also a magnetic drum machine, very much like the G15. Only one of those was built, for the British Motor Corporation.

This article is based on a paper presented at the Society’s all day NPL seminar held at the Museum on 20 May 1993.

The coming of mix and match hardware

Robin Shirley

The principle of mix-and-match hardware was introduced by the open architecture of the original Altair bus. It then became a standard that set the tone for the personal computer movement, and in other forms still continues with us today.

In late 1974, a series of construction articles on a rudimentary Intel 8008-based Mark 8 microcomputer appeared in *Radio-Electronics* magazine. This was the first time a computer had been put within the reach of anyone but a large company, and it aroused enormous interest.

Not wanting to be outdone, Les Solomon, the editor of *Popular Electronics*, commissioned Ed Roberts, the president of a small company called MITS in Albuquerque, New Mexico, to come up with a similar computer kit. Roberts decided to base it on Intel's new 8080A chip, and the Altair 8800 was born.

Reputedly the name Altair was suggested by Solomon's 12-year-old daughter, after the Enterprise's destination in the episode of Star Trek she had been watching. Roberts, being a sci-fi fan, liked it too.

The first Altair article appeared in the January 1975 issue of *Popular Electronics*. It had a bus based on a 100-way edge connector on which MITS had got a good surplus deal, and was accordingly called the Altair bus.

The resulting incidental but lavish availability of 100 bus lines led to the provision of a very rich environment of status and control signals, which was in turn to prove a spur to the designers of third party add-in boards.

Ed Roberts comes across at first sight as a somewhat unusual person to start an industrial revolution, though perhaps not in a deeper sense, since an unusual person is presumably just what's needed.

Seattle journalists James Wallace and Jim Erickson in their book "Hard Drive" (chronicling the rise of Bill Gates and Microsoft) describe Ed Roberts as "a hulking bear of a man". He was certainly physically big (six foot four and 21 stone), and he had a forceful and often overbearing manner (it's reported that only the hyperactive 19-year-old Gates, slightly built but with the confidence of his moneyed background, refused to back down to him, and that they had ferocious rows).

Roberts also had enormous energy and a powerful appetite for new

knowledge, new ideas and new activities.

Looking at his role in the S-100 and PC story, an equally important side to his character was his position as someone who, though within the electronic engineering industry, was so in a sense more as an amateur than a professional. His real ambition was to go to medical school and become a country doctor. That was exactly what he was to do in 1977 when he sold his business in its entirety to Pertec.

The infant personal computer industry was the creation as much of electronic hobbyists as of professionals. Roberts, despite owning a manufacturing company, was in close touch with the hobbyist community and felt at home among them.

His company MITS had begun as a genuine garage operation — it actually started operations from Roberts' garage in Albuquerque when he left the US Air Force. Initially he sold mail-order model rocket equipment and transmitters for radio-controlled model planes, and indeed the initials MITS originally stood for Model Instrumentation and Telemetry Systems.

Model would get changed later to *Micro* in the sort of retrospective promotion that often happened in the early industry — as for example with CP/M, Gary Kildall's ubiquitous 8080 disc operating system, known everywhere by its initials, behind which its original name was quietly updated to something more impressive.

If you look inside an Altair 8800, what strikes you is the way it's been put together by basic prototyping methods, such as slotted angle strips, with no hint of modern production engineering technology — a huge contrast with its polished successors of only three years later. This flavour of clumsy, cut-and-try manufacture seems to pervade all the early MITS production.

But it's important to realise not only how strikingly crude the Altair was, but also how crude it could afford to be (or even needed to be) if it was to catch its moment.

It's often said that if a thing is worth doing it's worth doing well. This was not at all the right proverb for the infant personal computer industry, which was better advised to say that if a thing was worth doing it was worth doing fast, immediately and cheaply.

This is well illustrated by contrasting the Altair with another class of microcomputer which, though far more professionally constructed and actually available earlier, was not to be on the evolutionary branch that led to the modern PC. These were the microprocessor development systems

made by corporations like Intel and Motorola. They were well-designed, solidly-built, on the bulky side and intended mainly for engineering development work—but at prices that were over 10 times that of the basic Altair kit, which MITS originally advertised for \$397.

This astonishingly low price was only possible because Ed Roberts had succeeded in browbeating Intel into selling him 8080A chips in volume for \$75 each instead of the regular price of \$350—the sort of enterprise that was to be crucial to the success of the Altair and its bus design, and hence to the launch of the personal computer movement.

Its internals may have been crude, but to the average hobbyist the great thing about the Altair was that it seemed attainable, whereas an Intel MDS was as far out of reach as a PDP-11 or Data General Nova.

In 1974, with MITS near bankruptcy, Roberts had bet the company on the Altair project, securing a \$65,000 development loan by convincing his bankers that he could sell several hundred of the machines. He had underestimated the irresistible magic of the idea of owning one's own computer. Within a few weeks, over 4000 prepaid orders had poured in and MITS leapt from some \$300,000 in the red to nearly as much in the black. The personal computer—a name coined by Roberts in advertising the Altair—had arrived.

The original Altair was essentially a prototype and had many shortcomings, from a feeble power supply to somewhat flaky bus timing. It was replaced in due course by a revised production version, the Altair 8800b, which was somewhat better.

A number of improved clones started to appear, so that by August 1976 *Doctor Dobbs' Journal of Computer Callisthenics and Orthodontics* (DDJ) was calling its 100-way bus the Altair/IMSAI or Hobbyist Standard bus. Roger Mellen of the then small company Cromemco proposed the name Standard 100 bus, or S-100 for short, because it had 100 lines, and this was the name that stuck.

The S-100 bus had most of the faults and virtues of unplanned industry standards. It had been designed in a hurry, wasn't optimised against crosstalk, and leant rather too much on the peculiarities of a particular processor, the 8080. On the other hand, it could be made to work reliably and was good enough. It quickly became the de facto standard.

The S-100 bus mirrored the architecture of early microprocessors, first in that it contained an 8-bit data bus. Actually, given the abundance of lines available, it had separate input and output data buses, which later by

multiplexing were to provide a simple extension route to a 16-bit data bus. It also had a 16-bit address bus, and, from the professional standpoint of electronic design engineers, who tend to be minimalists, an unnecessarily rich collection of status and control lines, and clock signals.

Another important and I think prudent choice by MITS, which at the time reflected its need to get something workable out of the door quickly despite its lack of experience with computer systems, was the design decision to distribute unregulated rather than regulated power.

Thus the S-100 bus provides most of the voltages likely to be needed by chips, in the form of unregulated DC rails of nominally +8V, +/-16V. By unregulated, I mean full-wave rectified and smoothed by a single electrolytic filter capacitor (a huge 180,000 μ F affair the size of a baked-bean can, in the case of the main +8V power rail, which could deliver getting on for 30 amps in some machines — enough for light welding).

A side-effect of the presence of this size of filter capacitor is that enough charge is stored to let the system ride out mains dropouts shorter than about half a second. I have often seen a Horizon or similar machine sail blithely through a dropout that momentarily dims the room lights, while modern PC/AT-clones in the same room reboot or hang.

Since bus line 13 provides a power-failure warning signal which is guaranteed to give at least 50ms for action before the voltage rails go out of spec, it is also possible to implement a form of power-fail auto-restart on S-100 systems, or at least a controlled shutdown.

The unregulated power rails might retain one or two volts of AC ripple, so they are not suitable for feeding directly to ICs. Hence it was left to each card to be responsible for providing its own on-board regulation for the voltages it needed, typically with basic fixed voltage regulator chips like 7805's, 79L05's, 78L12's and 79L12's, giving +/-5V and +/-12V respectively.

These work simply and effectively (though inefficiently, dumping the excess as heat), generally with a smallish tantalum capacitor in parallel, which, if the design margins were cut too fine, would blow safely but alarmingly on occasion like an exotic fuse, producing a brief but impressive cloud of acrid smoke.

This probably wouldn't have seemed the most natural scheme to a professional circuit designer, who'd tend to regulate power centrally with more sophisticated switching-mode circuits, as in the IBM PC, but it had desirable characteristics for machines that were to be used and upgraded

by many different people in a great variety of environments.

The chief benefit of on-board regulation was the way it made mixed-vendor systems more docile, by stopping interference from propagating from board to board along the power rails.

It also helped to get working the relatively long buses with many slots (18 on the Altair, 21 on the Cromemco Z2), needed because of the limited functionality of the early boards—for example only 4Kb on each MITS memory board—so that a lot of boards could be needed in one system. It was these 4Kb memory boards in particular that were to prove a source of contention and encourage other manufacturers to challenge the MITS monopoly.

The basic \$397 Altair kit did not include any input-output other than the front panel display and switches, and also only 256 bytes of RAM. And of course no software. This was just enough for a hobbyist to get the thrill of entering a program via the front panel switches and seeing his own handiwork come to life and display a test pattern on the LEDs (if he was lucky). To do any useful work, more was needed—especially more memory, preferably enough to run a high-level language.

This was where Bill Gates and his partner Paul Allen came in, with what was to become Microsoft BASIC, developed semi-clandestinely using an 8080 simulator on the PDP-10 at Harvard, where Gates was then a second year student. Allen joined MITS as software director in spring 1975, and Gates a few months later, dropping out of his studies to the dismay of his mother, although Gates remained a freelance and was never formally a MITS employee.

By the early summer of 1975, the first 4Kb (kilobytes) paper tape version of Microsoft BASIC was shipping (made practicable by the ready availability of government surplus ASR 33 teletypes), soon followed by an 8Kb version, then an extended BASIC requiring 12-16Kb of RAM (programmers have never had any difficulty overflowing the available memory).

From MITS' point of view, the purpose of Microsoft BASIC (which it marketed as Altair BASIC) was to sell hardware, principally memory boards, and its contract with Microsoft was drawn up with that in mind.

Dynamic RAM (DRAM) chips store data as charges in tightly-packed rows of capacitors, which have to be refreshed every few milliseconds before they leak away, whereas static RAM uses transistorised flip-flop switches that need more power and space per bit, but are simple to implement. At any given time, DRAM always offers more bits per chip and per dollar

than SRAM, but with less speed and more problems.

Unfortunately, the early DRAM chips had particularly finicky timing and voltage constraints and were difficult to design for—so much so that a number of memory board manufacturers, notably Bill Godbout, were to hold out for years against DRAM, sticking to the more expensive but more forgiving SRAM. At MITS, however, Ed Roberts predictably picked the low-cost DRAM option, but found that he had bitten off more than he could chew—the MITS 4Kb DRAM boards were poorly designed and seldom worked as advertised.

Meanwhile, other companies started to ship their own RAM boards which did work. However, in order to get hold of Altair BASIC, people were having to buy the near-useless MITS memory boards. Even then, they weren't getting the copies of BASIC that they had ordered and paid for, because of shipping delays at MITS due to the memory board problems. How this led to the first piracy of personal computer software and the split up of Microsoft from MITS is a fascinating story that you can read in "Hard Drive".

The time window in which MITS could successfully exploit its vision in creating the personal computer market was to be a relatively brief one. A California company called IMS Associates Inc (IMSAI) looked at the Altair and decided it could do better. The IMSAI 8080, the first Altair-compatible machine, appeared in late 1975.

The Altair with its front-panel of metal toggle switches and LEDs looked like a proper computer, specifically a Data General Nova. The IMSAI 8080 went one better, with a row of coloured-plastic flip-switches that looked just like a PDP-11. The overall quality of the IMSAI was also much higher, so that here, at last, was the basis of a reasonably reliable small business microcomputer system.

It was soon followed by the handsome but expensive Processor Technology Sol, with its smart blue livery, built-in keyboard, low-profile design and horizontally mounted bus slots, and the Polymorphic Systems Poly 88.

Floppy disc systems were appearing too, at first 8-inch drives, holding a nominal 250Kb per diskette, and later the 5.25-inch Shugart SA400 mini-floppy format, which in single-sided single density form stored about 175Kb. Among the add-on disc system vendors was a California company called North Star, whose products were bundled with a free but spartan disc operating system and an excellent BCD-based BASIC interpreter, and

whose blue-painted drive cabinets soon became a common sight.

North Star also made a hardware floating point board designed around the 74LS181 4-bit ALU, which too was supported by versions of North Star BASIC. Just as would occur a decade later with add-on PC board makers, it was soon to use this experience as a springboard into producing complete systems.

The period 1976-77 saw the arrival of a host of high-quality second-generation designs based on the 4MHz Zilog Z80A, which ushered in the golden age of S-100 systems.

As the personal computer movement gained in confidence and maturity, it was realised that handswitches were no longer necessary. The new systems instead jumped on reset to a bootstrap PROM on either the CPU board or the disc controller, which then booted up the operating system from a floppy disc. There followed a host of classic designs on this pattern, such as the Cromemco Z2, North Star Horizon and Vector Graphics MZ.

By 1979 the future was already fixed, although in terms of market share and value, personal microcomputers made up a cloud on the horizon no bigger than a man's hand.

Anyone who stood back at that time and studied the evolutionary ecology, as it were, of computers, could see at once that in the long run microcomputers were bound to replace the traditional mainframe and mini-computer architectures. It was simply a matter of their evolution rate and generation times.

Whereas it took seven to 10 years to produce a new generation of mainframes, and some three to five years for minis, microcomputers evolved at a simply phenomenal rate—initially at better than a generation a year, levelling out at one to two years. They simply outbred their rivals.

If we label the rough and ready Altair 8800 of early 1975 as the zeroth generation, and the more reliable and better engineered IMSAI 8080, Sol, Poly 88 and Altair 8800b of 1975-76 as first generation, then polished second-generation Z80A-based machines that rivalled low-end mini-computers, like the Cromemco Z2 and North Star Horizon, were already appearing by 1976-77.

By extrapolating this graph, driven as it was by the inexorable progress of Moore's law concerning the doubling time of IC circuit density, one could even then predict a crossing point in the early 1990s when microcomputers would start to pull ahead of all other types, as has now happened (though of course often in disguise, like DEC Alpha minis and

parallel supercomputers).

After considerable effort over several years, through the efforts of people like George Morrow of Thinker Toys, the extension of the S-100 bus definition to a new IEEE 696 standard was achieved, producing an upwards-compatible extension from an 8-bit to a 16-bit data bus, and from a 16-bit to a 24-bit address bus, without requiring any changes in pin layout or in most cases to motherboard design. The draft standard was published in July 1979 and the final version approved in June 1981.

According to Sol Libes and Mark Garetz in their classic 1981 book “Interfacing to S-100/IEEE 696 microcomputers”, there were then some 200,000 S-100 systems in operation. It’s interesting to bear this figure in mind (considered enormous at the time) when looking at what was soon to follow. There were also reported to be nearly 100 different manufacturers offering about 400 different plug-in S-100 boards.

Emphasising the aim of processor independence for the IEEE 696 bus, Libes and Garetz list eight different 8-bit CPU boards that were available for it (8080A, 8085, Z80A, 2650, 6502, 6800, 6802 and 6809, although in reality the first three are variants on the same architecture and made up nearly all the total). They also listed seven 16-bit CPU boards (9900, LSI-11-like, 8086, 8088, Z8000, 68000 and Pascal Microengine).

In general, upwards compatibility from earlier hardware was very good, although conflicts with the new standard inevitably occurred with some of the early designs, where undefined lines in the original Altair specification had been used for proprietary purposes. In practice this was relatively unimportant: few such machines had been built and they were unlikely to be upgraded with newer processors.

Conflicts also arose in some later designs which had tied to ground, for reasons of improved noise immunity, previously unused lines which became defined in the new standard. An example was the North Star Horizon motherboard, where the unused bus line 61 had been tied inaccessibly to ground beneath each edge connector. Unfortunately, when the IEEE 696 standard came along, this previously unimportant line became A20 in the extended 24-bit address bus, which meant that Horizons required significant modifications before they could be upgraded to new 16-bit processors.

In practice this too proved unimportant, since Horizons, like most other machines, tended to carry on doing the tasks they had been bought for rather than get involved in major upgrades. This is a story with a moral for those today who pay extra for ‘upgradable’ and ‘future-proof’ computers:

the lesson of the short but breakneck history of microcomputers is that no such animal has ever existed, or is ever likely to.

In fact none of the 16/24-bit extensions nor indeed the IEEE 696 standard itself were in the end very significant, because the whole episode proved to be another example of the sailing ship effect. By a classic piece of technological irony, IEEE 696 was finalised just as the whole thriving S-100 bus scene was about to be overshadowed by the introduction of the IBM PC and its successors. A mere four years later, with the coming of the PC/AT in 1985, it was fading into irrelevance. The 200,000 S-100 systems of 1981 had been overwhelmed by the millions of IBM PCs and clones.

We have a similar situation today, in which the history of the IEEE 696 story has repeated itself (also after some five or six years) in the case of the 32-bit EISA extension of the PC/AT bus. This too is coming into widespread availability just as it is about to be superseded, first by the VL (VESA local) bus, and probably soon also by Intel's new 64-bit PCI (Peripheral Component Interconnect) bus standard.

This article is an edited version of the talk given by the author to the Society at the Science Museum on 17 June 1993.

New contact point

Readers wishing to contact the Secretary are reminded that he is now running the secretariat from his home, and can no longer be contacted at the Science Museum.

The new secretarial telephone number is 0234 822788. Letters should be addressed to Tony Sale, Secretary, Computer Conservation Society, 15 Northampton Road, Bromham, Beds MK43 8QB.

NRDC's role in the early British computer industry

John Crawley

The NRDC was created as the end product of a long period of policy debate, which started before the first world war. As a result of that war, there was a lot of discussion about what should be done to try and stem what was perceived as the decline of industry and science in the UK.

Nothing came of this, but the subject came to the surface again early in the last war. Then far-sighted people like Blackett saw there was a need to ensure that the results coming out of Government-supported research for defence purposes could ultimately be used for the benefit of UK industry.

A debate ensued. The Board of Trade, the Treasury and other departments were all arguing about it with their own views on what should be done. By about 1943 some solid ideas started to emerge, and by the end of the war these were taking shape. Under the aegis of Harold Wilson and Stafford Cripps a Parliamentary bill was drafted and debated. This resulted in the Development of Inventions Act of 1948.

The concept of the Corporation at that time was based on the primacy of patentable or patented inventions. The rather simplistic view was that if you had patents, you could get industry to take licences and they would use the inventions. The process would go on naturally as they would produce products and sell them and this was all good for industry.

The discussions were also complicated by the fact that the people involved were primarily concerned with inventions that arose from research conducted in publicly supported places, such as research establishments and universities. The principle applied that work done under publicly supported contracts would result in the industrial property rights belonging to the Crown, so that the Government departments responsible would then have the task of trying to do something about these rights. So that emphasised again the concentration on handling patents.

A further constraint then was that a general principle of fairness that had to be applied. Contractors couldn't be given any unfair advantage: rights couldn't be left with them, and they couldn't be chosen specifically. Things had to go through an open tendering processes.

The Act set up an independent executive body, the National Research Development Corporation: not a Civil Service organisation, although it owed allegiance to the Board of Trade. In order to perform its functions

it had powers to borrow up to £5,000,000 from the Board of Trade, and it was required to repay this money, I think, within five years with interest. It sounds nonsense now but then it seemed reasonable. Subsequently there were later amending acts which increased the borrowing ceiling and extended the repayment period indefinitely.

The functions of the Corporation as defined by the Act were twofold. First, securing, where the public interest required, the development or exploitation of inventions resulting from research paid for with public funds; and also of any other invention which NRDC itself considered was not being sufficiently developed or exploited in the public interest.

(That meant that the Corporation had to do something to secure development and exploitation of results arising from public research. If something else from other sources came into its view and NRDC felt that something ought to be done, it had powers to act. In many cases it did, most notably with the hovercraft.)

The second function was to deal with inventions resulting from public research and with inventions from other sources. This amounted to looking after the licensing of inventions, including private inventions.

I want to go back a bit in time to before the Corporation was created. During the latter part of the war years there was a substantial drive, primarily by the Ministry of Supply, to follow up the work that was going on in the various service research establishments. This was to see if there were any inventions and new technologies emerging which could be protected by patent action and would subsequently would be of benefit to industry.

I got involved in this myself towards the end of the war—I was temporarily attached to the Ministry of Supply at that time. We were following up work which had been started at the Telecommunications Research Establishment (TRE) by F C Williams and Tom Kilburn and was subsequently continued at Manchester University. This work was concerned with immediate access digital storage (the Williams tube cathode ray tube store) which at that time represented a most important advance in the electronic techniques then available for realising electronic digital computers, which had begun to excite the informed people at that time.

After the war I went back to my own employment on the examining staff of the Patent Office. But I hadn't been back there very long when I was borrowed back again by the Ministry of Supply. By that time the NRDC was on the verge of creation and they wanted to get on with preparing the patent holdings to be passed over to the NRDC.

The NRDC was formally set up on 28th June 1949 with Lord Halsbury as managing director. A Board was brought together which comprised quite a number of prestigious scientific establishment people, like Patrick Blackett and Sir Henry Tizard.

A small nucleus of staff was brought together, including Dennis Hennessey, a patent agent involved in the Ministry of Supply patenting drive, and myself.

It was clear from the outset that the exploitation of what was by then a quite substantial block of patent rights or patent applications would be a major task for the Corporation. We thought initially in terms of licensing and we hoped that the techniques would be exploited by appropriate sectors of industry. But by whom, we didn't know.

Significant events which were profoundly to affect the future were taking place before the Corporation was formally in existence. IBM heard about the immediate access Williams storage tube work and got in touch with him, saying they wanted to know all about it. Although the company was not seriously at that time contemplating becoming commercially involved in computers as part of its mainstream business, it was building some machines.

The result was that Halsbury and Williams went to the States, I think towards the end of 1949, and fixed up an option agreement with IBM for a licence. Early the following year that was converted into a proper licensing arrangement, after IBM had taken the decision to get on with the Defence Calculator and the 700 series machines.

That licensing arrangement was the first business thing the Corporation did, and it turned out to be a very satisfactory one. But we were not set up just to earn money from the Americans: we were supposed to be doing something in the national industrial interest.

From the beginning Lord Halsbury was convinced that the exploitation of the rights held by the Corporation would be related to business and data processing applications of computers, and required the encouragement of the building up of an effective computer industry in the UK, or positive assistance to ensure such industry creation. He wanted to do just that though some people were arguing that it was not a proper role for the Corporation to undertake.

Lord Halsbury said the only way to do it was to cause one of the major electronic companies to get together with a business machine organisation, such as Powers-Samas or British Tab, to develop and build something

and get it on the market. So we tried hard to force some kind of joint venture activity between an electronics firm and the punched card business machine companies, Powers-Samas or British Tab.

This was a painful period that lasted 18 months to two years when we had a succession of working parties and study groups and meetings with industry. It all came to nothing, and it is clear to me in retrospect that companies concerned were in general suspicious of the NRDC and of Government interference.

Neither did any of them really believe that there was a commercial future for computers. Some of them were arrogant enough to say, "If it turns out that there is a need for these things, we are quite satisfied that our company could deal with it ourselves without your help".

The first record I have of Halsbury enunciating this philosophy of forcing a juncture between companies was made—it's in the records—at the first Board meeting of NRDC in July 1949. This phase of exhortation, persuasion, pressure or encouragement was something that Halsbury later referred to as "time spent trying to push mules uphill".

That period lasted until early 1951, by which time discussions were starting between the Corporation and Ferranti and Elliott Bros about supporting some computer developments. Ferranti was already involved in the computer business because it had a contract from the Ministry of Supply to provide the engineering support for the team at Manchester University.

This was a far-sighted move by Ben Lockspeiser, then the Chief Scientist. He gave an open ended contract to Ferranti to provide these services to Manchester University without any question of tendering.

As Ferranti was already in the business of making machines, we were well aware of what they were doing. We were not so well informed then of what was going on at Elliott Bros, although we had learnt something of it as a result of the meetings over the previous 18 months. We also had contact with Lyons but they seemed to be quite self-contained and didn't seem to want anything to do with us.

Another impression I have is that the punched card people had not grasped the idea that it was the potential for business applications that motivated us, and thought that we were inviting them to do something to make powerful machines for scientific computation.

Towards the end of 1950 Halsbury was in the States again with IBM. He found out all that was going on there and was very impressed. He

somehow got the idea — perhaps it was true — that there was a government requirement for 10 machines in the States which were wanted during 1952. There appeared to be no chance of getting the machines in time from any source in the States.

Halsbury thought, “Why don’t we supply these?”. There was a cable bombardment across the Atlantic trying to persuade Hennessey and the rest of us back in Tilney Street to try and get one or other of the firms developing and making something to deliver by 1952.

This also came to nothing. Elliott did say that they could perhaps produce something, but not what was apparently wanted, which was something like the Ferranti/Manchester University machine. They could make and deliver in 1952 10 machines according to their own existing design. That came to nothing either.

But this exploration resulted in us having much more exposure to Elliott Bros. We saw their current work, and there was also reference to some use of Williams storage tubes which interested us considerably. We were also impressed with what we saw of their construction techniques.

The result was that some time in 1951 we gave Elliott Bros a study contract to consider whether their constructional technique could be applied to something like a re-engineered version of the Manchester University machine. That led to a number of reports and proposals for a machine.

In the end a proposal emerged for the development of a packaged computer. The suggestion was that we should pay for the development and construction of a prototype — this turned out to be the 401. There was an outline specification suggested by them and as usual in these things nobody could ever agree about anything.

We put the problem of deciding whether it was a suitable design for us to support up to the Brunt committee. That was the advisory committee on high speed calculating machines which had been set up by Ben Lockspeiser when he transferred to the Department of Scientific and Industrial Research, to monitor and look after the work which he had started with his contract with Ferranti. It was chaired by Sir David Brunt, and its members included Professors Hartree and Wilkes.

The result was the Brunt committee after brief deliberation said yes, they thought it was all right. So we went ahead and gave Elliotts the contract. We got the approval (we had to go to the Board of Trade for these things) by April 1952, but by that time work was already going ahead and was quite well advanced.

We extended the contract to cover the development and construction of a prototype. I think the estimated cost was about £30,000 for the whole job, and when it came to the end and the 401 went off to Cambridge and then to Rothamstead, I don't think the whole project had cost us much more than £50,000.

It was very interesting that in the arguments we had to get approval for it, and even internally, there was a suggestion that by going to Elliott Bros and giving them a contract we were setting up a competitor to Ferranti. It's ironic that if it had been the other way round people would have been chiding us for giving people an unfair advantage.

In parallel, while this was going on Halsbury had been trying to get the Ferranti people to develop a commercial machine. That ended up with a fairly straightforward and quickly settled contract whereby we undertook to pay for the construction of four, subsequently six, copies of the Manchester machine.

This became the Ferranti Mk 1*. The agreement was that they would make six machines which we would pay for, and they would resell them as our agents at prices which we would mutually agree. This would give Ferranti their costs plus profit at 7.5%. If we set the selling price right it would also bring back to NRDC what we had spent with some profit. And it did; it worked very well.

That arrangement with Ferranti was running before the 401 project with Elliotts was set up, sometime towards the end of 1951. The success of the Ferranti Mk 1* project made us feel that it represented a good method of going ahead to provide support. We had no reason to believe that anything could go wrong, but subsequently it did, when we later turned to Ferranti to carry on developments based on the Elliott work.

John Coales, the Director of the Elliott Brothers Research Laboratories (and the main contact between the company and NRDC), decided in around March 1952 to leave the company. Whatever his reasons, his departure had a disturbing and damaging effect on morale.

Later that year cutbacks in Admiralty support threatened staff cutbacks at Elliotts. Fortunately these did not occur, but morale had been further damaged: people began to look for other jobs, and some did actually leave the company.

The situation as presented to NRDC appeared much more serious than it eventually turned out to be. It raised questions about the prospects for further development and exploitation of the technology by the company,

even if it completed construction of the prototype 401.

Halsbury and the members of what had by that time been created as a subcommittee of the main board of the Corporation (the computer subcommittee; I, for my sins, was the secretary of that) were convinced that Elliotts computer R&D effort was just going to disintegrate, and what we had hoped would be continuing further developments and manufacture and commercialisation of something based on the 401 couldn't happen. Why did we form this view so strongly? I don't really understand now. We were forced to the conclusion that we had to set up some rescue operation, which we did.

What subsequently emerged was that we were quite wrong in our forecast of what was going to happen at Elliott Bros; they went on from strength to strength and did very well. And as far as the 401 was concerned, work continued satisfactorily, and it was exhibited at the Physical Society Exhibition in April 1953.

I think it's a great pity we didn't continue to support Elliott Bros as they might have been a much better base for a larger R&D effort. This might have led earlier to the next generation of computers.

At any rate, the NRDC concluded that further development and commercialisation of the 401 was unlikely, so we tentatively explored the possibility of other companies being induced to carry on the work, with the help of what could be kept together of the 401 team. Ultimately this resulted in the involvement, once again, of Ferranti.

Initially we did not seriously consider Ferranti, as we felt the company was fully occupied with the Mark 1*. But by early 1953, the company had absorbed some of the staff who had left Elliotts. In September of that year Ferranti offered WS Elliott, Head of the Computer Division at Elliotts, an appointment, and this prompted NRDC to consider seriously making 401 technology available to Ferranti.

By the end of 1953 Ferranti had absorbed more former Elliott Bros staff, including Charles Owen, Hugh Devonald and George Felton. The company then proposed to NRDC that it should underwrite the development and production of computers based on the 401 package technology but with many other features. An outline specification for what was called the FPC1 was produced.

This produced a strong difference of opinion between Ferranti and Christopher Strachey on behalf of the NRDC, who produced his own alternative specification. The result was a compromise specification, and on

this basis the project to manufacture what became known as the Pegasus eventually got under way.

The contract governing the project was for development, the production of a prototype, and then the manufacture of nine production machines, to be purchased by NRDC for resale by Ferranti — substantially on the terms used for the Mark I* contract. A limit on cost to NRDC of £220,000 (later £250,000) was included, all rights would vest in NRDC and Ferranti could not sell machines of the same design produced at their own cost before the NRDC machines were sold.

These contracts always tailed behind the actual work. Before it was formally signed, Sir Vincent de Ferranti refused to accept the terms, objecting to any requirement which might affect Ferranti's freedom of action. The work by this time was so advanced we couldn't possibly stop it and take it elsewhere, so we had to agree a compromise contract which removed the ceiling on expenditure. That didn't worry us much at the time, because we thought the figures could be relied on. The new contract also did some other things about rights and we originally had a veto over design changes to safeguard Strachey's position.

The revised contract was eventually signed, and then we discovered that the costs were escalating. There had been some overheads which had been lost, and we were stuck with costs which eventually went to close on £500,000, instead of the £250,000 or so which we'd originally thought. By that time orders had been taken for at least eight of the machines at prices pitched on the basis of the earlier costings, so we lost a lot of money over it.

Despite these financial difficulties, the 401 initiative must be considered a great success. Two streams of computer development and manufacture emerged from it: the Elliott 400 series, which in its 405 version achieved something like the target Lord Halsbury was aiming for; and the Ferranti Pegasus and its derivatives.

This article is an abridged version of a talk given by the author as part of the Elliott/Pegasus all-day seminar at the Science Museum on 21 May 1992.

Working Party Reports

Elliott 401

Chris Burton, Chairman

Conservation of the major sub-assemblies of the computer has continued to forge ahead, to an extent overtaking the ability of the Working Party to keep up with pre-commissioning work. Accordingly, we decided to revert to an earlier and simpler plan, abandoning the idea of using temporary power supplies to power the logic while the real supplies were being conserved. Now the latter will be conserved first — indeed, the work is almost complete — and commissioning of the whole system will start with the power system.

Elucidation of the logic diagrams and order code is turning out to be very tedious and drawn out, but Peter Holland and Maurice Hill valiantly and patiently move the work forward. Maurice has created a bit-level simulator program of parts of the logic to help verify how it should work. Recently we were very pleased to receive some information from a one-time user of the 401, giving sample program fragments.

Very satisfactory progress has been made on recovery of information from the tracks on the drum. A somewhat Heath Robinson lash-up on the bench allows us to capture the analogue track signal as a digitised waveform in a file on a PC. We then decode the signal using a purpose-built analysis program, together with some manual, or at least mental, techniques. Peter elucidated Track 0, the initial orders track, and later was delighted to find independent evidence of the correctness of some of the data. The hard task next is to work out the semantics of the routines. Len Hewitt will continue to capture good track files until we have a complete set, and we can then do analysis of the signals at leisure at any time in the future.

The next major task is to choose the moment to move to Blythe House. It will not be worth starting the system commissioning before the move. We look forward to setting up our new base in the Elliott room there.

Pegasus

Chris Burton, Acting Chairman

Very little activity has taken place on the machine since the last report, partly because of the less flexible arrangements for access in the absence of Tony Sale, and partly because of the usual holiday season lull. We had one request to run and demonstrate the Pegasus for the BCS Branches Board meeting at the Science Museum, when 40 visitors saw the machine operating. It remains adequately reliable for demonstrations, requiring half an hour to warm up before it reaches stability. The main task for Working Party members will be to try to restore the margins as they were at the beginning of the year in an attempt to eliminate the warm up period.

We have been discussing appropriate ways of distributing replica Pegasus documentation with the simulator. We have not yet reached a conclusion, as we try to find a path between the aesthetically desirable and the economically practical.

There is no commitment to move the Pegasus to a new site yet, so we continue to plan on the basis of its remaining in the Old Canteen, whatever use is made of that building. Operations are likely to become more difficult as the focus for Society activities moves to Blythe House: for example, we all share one set of tools and test equipment at the moment.

We have made two visits to the Pegasus 1 at the Manchester Museum of Science and Industry, and have submitted a detailed report on its condition to the curator, Dr Jenny Wetton. Members have contributed to the attempt to ascertain the serial number of the machine, which is thought to be either number 1 or, more likely, number 6. Now that the North West Group of the Society has been inaugurated, a local working party will be set up to handle the CCS involvement with the machine.

S-100 bus

Robin Shirley, Chairman

I went to Manchester in July to acquire for the collection an interesting home-constructed S-100 system, based on a kit version of the UK-designed Transam Tuscan microcomputer, and donated by its owner, Mario Wolczko (who was about to leave the Computer Science Department at Manchester for a post with Sun Microsystems in Silicon Valley).

Apart from providing an example of Tuscan-based equipment, Mario's system is notable for the outstanding completeness of its documentation

and provenance, which includes all the original invoices, correspondence and manuals, and a set of exemplary bound logs with schematics, etc, for all the hardware and software modifications which Mario made since acquiring the original Tuscan kit around 1981.

He has also loaned a set of slides for us to copy, recording various stages in the system's construction. The system includes a full front panel designed by Mario, with switches and bus status LEDs, and an external 8-inch floppy drive unit, and is housed in twin wood-varnished cabinets.

Elliott 803

John Sinclair, Chairman

By the time this issue of *Resurrection* reaches you, the Elliott 803 should be installed in Blythe House. We carefully dismantled the system in the Old Canteen in mid-October, under the watchful eye of Tony Sale and his video camera, so at the time of writing it consists of a set of packing cases. The plan is to transport the system to its new home on Tuesday 9 November.

After that we hope to be able to start restoration work again. A necessary first step will be to organise a new tool kit: the ones we have been using up to now have remained in the Old Canteen, as they will be needed for work on the Elliott 401 and the Pegasus.

DEC

Adrian Johnstone, Chairman

As with most of the other working parties, restoration work has been suspended while we have been waiting for the move to Blythe House. This is scheduled at the time of writing for the second week in November. After that we are looking forward to resuming activity.

One development during our period of enforced idleness is that we have been offered a machine by the Atomic Weapons Research Establishment at Aldermaston. It was previously used in the nuclear weapons testing program.

Letters to the Editor

Dear Mr Enticknap,

In issue 6 Cecil Marks refers to the HEC range of early computers produced by the British Tabulating Machine Company. Having been involved with their development I can fill in a few gaps.

There were four versions. HEC1 was an experimental one designed by Andrew Booth. This was followed by a prototype, HEC2, designed by Richard Bird. A modified version of this, designated HEC2M, became the production version, of which, speaking from memory, five were sold. The price was £12,000. The first was delivered to the GEC Research Laboratories at Wembley early in 1955 and the second to the Esso Refinery at Fawley in July of that year. Both were used successfully for a wide range of engineering calculations, but were not really suitable for routine office tasks. As Mr Marks says, the first version to be designed for general data processing was the HEC4, later to be called the 1201.

Unfortunately, to the best of my knowledge none of these machines remain.

Yours sincerely,

Brian Dagnall

Lymington, Hampshire

29 July 1993

Dear Mr Enticknap,

In writing about “The Design of Pegasus” in issue number 7 (autumn 1993) Ian Merry identified Charles Owen, Christopher Strachey, Brian Maudsley and himself as the major players. One notable, and in my view regrettable, omission is CH (Hugh) Devonald.

As Ian points out, Owen and Strachey “aimed to build all of the complex control functions without recourse to special purpose circuits”. This approach meant that, once the logical functions and the interconnection rules of Owen’s standard circuits had been defined, the overall logic design could proceed as a separate quasi-mathematical activity whereby the requirements of the machine’s architectural design were analysed and a logic structure to meet them described in appropriate symbolism, which could then be expressed in the form of lists of packages and their interconnections. This logic design was the work of Hugh Devonald and his

team: and Ian Merry was not alone in finding that “the description of the various control cycles was and remains baffling”.

As Ian says, some of the leading members of the Pegasus team, notably Elliott, Owen and he himself, left Ferranti in 1956 shortly after the first Pegasus was completed (it was built not in a factory but in the computer room on the first floor of 21 Portland Place) and before the first factory-made machine was delivered. Hugh Devonald then became a key figure in the continuation of the work which included the development of a magnetic tape system, associated card-to-tape and tape-to-card or -line printer converters (also usable with the Mercury computer) and the stretched Pegasus II. Even more significant was his continuation to completion of the Perseus project which without him could not conceivably have survived the 1956 departures.

Perseus was built largely from Pegasus packages supplemented by long torsional-mode delay lines. It was a remarkable machine with a word length of 72 bits holding either 12 6-bit binary coded characters or three 24-bit instructions. It had a self-checking mixed-radix arithmetic unit and was designed for commercial data processing. The main computer comprised eight cabinets as against the three of a basic Pegasus. It says a lot for the quality of the design of Pegasus that the methodology and technology worked splendidly even when stretched to this extent. No prototype was required and the two production machines (delivered in 1959) performed very well for their users in Sweden and South Africa for many years. However, Perseus was stretching valve technology rather far, and the onset of transistor-based machines limited its selling life.

Yours sincerely,

MH Johnson

Oxford

6 October 1993

Dear Sir,

The interesting article in issue 7 by Doron Swade came, I think, to the right conclusion about preserving software, but missed some important points.

It helps to decide first what we are trying to achieve.

As to the purpose of a museum, there are two schools of thought which were well brought out in a paper by Gwen Bell (the current ACM presi-

dent) when speaking of the Boston Computer Museum at the recent Forum on the History of Computing¹.

There is the Educational Purpose, which the Boston museum does extremely well — lots of demos to delight the kiddies, who will indeed come away having learnt much about computers. There are also static exhibits of artefacts, set in appropriate contexts to illustrate their significance.

The second Purpose is the preservation of “worthy artefacts”, presumably to facilitate their study in the future. Here the policy of the Boston museum was not so convincing. For example, they had been offered the last working Multics system. Should they accept it? I asked whether they had the means to maintain it in working order and was told, “Oh no! That is not the purpose at all”. But to preserve it as a static exhibit would be pointless. The hardware itself is not especially interesting — it is the pioneering Multics operating system that is so important and, we should hope, of interest to future students.

To digress for a moment, let us ask *why* a serious student would wish to use a museum? The answer is to understand the past: in particular, to do so by trying to get inside the mindset of the original user of each object. So if the museum displays a Stone Age object, the archaeologist, having observed it, can (and indeed does) then try to make one for himself, thus replicating the technology of 20,000 years ago.

What of the technology of 200 years ago? You can examine a Watt steam engine (even a Babbage Difference Engine) and, by visual examination alone, determine how it worked and how it was manufactured.

The technology of 20 years ago? Imagine trying to understand, with a view to replication, the Pilot Model ACE (or even a well-preserved Pegasus). The mind boggles, but I suppose it could be done if one were desperate enough (the importance of preserving manuals springs to mind).

But what of the technology of two years ago? What use is it to preserve a Sparc chip? Just a piece of black plastic with leads sticking out. Even if you could determine its circuitry by destructive examination, no-one in future is going to rebuild the manufacturing technology necessary to duplicate it (one might just about manage the crude technology necessary to reproduce a Z80).

The moral is that the curator also needs to put himself into the right mindset, this time that of the student of 50, 100 even 1000 years hence, and then to choose his objects and methods of preservation accordingly.

¹Forum on the History of Computing, held at Boston, Ma, 20 April 1993

So let us now apply this principle to software. A set of diskettes on which Windows 1.0 was once recorded will make a pretty, but not a helpful, exhibit. If some future student wants to get into the mindset of today's programmers, he will want to study, and *use*, the preserved software. But software is just bit patterns — the medium on which it is recorded is totally irrelevant.

Currently, the best way to hold large amounts of archived software would be a computer with many gigabytes of disc and a proper tape back-up facility — a Unix box would do nicely. The one system can preserve software arising from all sorts of sources, and that once ran on all sorts of machines — and all nicely indexed and cross-referenced to boot. In 10 or 20 years time your Unix box will be obsolete — no problem! Just move the bit patterns over to whatever system is then in fashion. Bit patterns can be preserved indefinitely.

So what will the student of the future do with it? Two things:

1. Study it. Therefore it is most important to preserve the source code, if at all possible, a point regrettably not mentioned by Doron Swade.
2. Run it — not on the original hardware, but on emulators.

Already, I hear, we can emulate EDSAC I on a PC, and marvel again at the elegance and economy of David Wheeler's initial orders, and the fundamental principles of software engineering laid out in Wilkes, Wheeler and Gill, which our librarian must be sure to have preserved carefully. But in 20 years time the PC will also have followed the dodo. Again no problem. It will have been emulated on something else, and the EDSAC I emulator will run on that emulator, and probably still run faster than EDSAC I itself.

So that is something that could be done — and at not all that much expense either. Who should do it? The Curator? The Librarian? The Archivist? What difference does it make? It is just a matter of semantics!

Yours sincerely,

CH Lindsey

Cheadle

18 October 1993

Forthcoming Events

2 December 1993 Evening meeting

The subject will be the Stantec Zebra - speaker yet to be finalised.

24 February 1994 Half day meeting

Debate on why the British computer industry did not capitalise on the country's lead in computing, starting 2.00 pm at the Science Museum (*subject to confirmation*).

19 May 1994 All day seminar

The design and development of the IBM 360 series, starting 11.00 am at the Science Museum (*subject to confirmation*).

All evening meetings take place in the Science Museum Lecture Theatre and start at 5.30 pm.

Resurrection is the bulletin of the Computer Conservation Society and is distributed free to members. Additional copies are £3.00 each, or £10.00 for a subscription covering four issues.

Editor – Nicholas Enticknap

Typesetting – Nicholas Enticknap

Typesetting design – Adrian Johnstone

Cover design – Tony Sale

Printed by the British Computer Society

©Computer Conservation Society

Committee of the Society

Chairman **Graham Morris FBCS**, 43 Pewley Hill, Guildford, Surrey GU1 3SW. Tel: 0483 66933

Secretary **Tony Sale FBCS** 15 Northampton Road, Bromham, Beds MK43 8QB. Tel: 0234 822788.

Treasurer **Dan Hayton**, 31 High Street, Farnborough Village, Orpington, Kent BR6 7BQ. Tel: 0689 852186

Science Museum representative **Doron Swade**, Curator of Computing, The Science Museum, Exhibition Road, London SW7 2DD. Tel: 071-938 8106

Chairman, Elliott 803 Working Party **John Sinclair**, 9 Plummers Lane, Haynes, Bedford MK45 3PL. Tel: 0234 381 403

Chairman, Elliott 401 and Pegasus Working Parties **Chris Burton CEng, FIEE, FBCS**, Wern Ddu Fach, Llansilin, Oswestry, Shropshire SY10 9BN. Tel: 0691 70274

Chairman, DEC Working Party **Dr Adrian Johnstone CEng, MIEE, MBCS**, Dept of Computer Science, Royal Holloway and Bedford New College, Egham, Surrey TW20 0EX. Tel: 0784 443425

Chairman, S100 bus Working Party **Robin Shirley**, 41 Guildford Park Avenue, Guildford, Surrey GU2 5NL. Tel: 0483 65220

Editor, Resurrection **Nicholas Enticknap**, 4 Thornton Court, Grand Drive, Raynes Park SW20 9HJ. Tel: 081-540 5952 Fax: 081-715 0484

Archivist **Harold Gearing FBCS**, 14 Craft Way, Steeple Morden, Royston, Herts SG8 0PF. Tel 0763 852567

Dr Martin Campbell-Kelly, Department of Computer Science, University of Warwick, Coventry CV4 7AL. Tel: 0203 523196

George Davis CEng FBCS, 25 Manor Way, Purley, Surrey CR8 3BL. Tel: 081- 660 5581

Professor Sandy Douglas CBE FBCS, 9 Woodside Avenue, Walton-on-Thames, Surrey KT12 5LQ. Tel: 0932 224923

Chris Hipwell, Bretts, The Green, Newick, East Sussex BN8 4LA. Tel: 082572 2567

Dr Roger Johnson FBCS 9 Stanhope Way, Riverhead, Sevenoaks, Kent TN13 2DZ. Tel: 071-631 6388

Ewart Willey FBCS 4 Sebastian Avenue, Shenfield, Brentwood, Essex CM15 8PN. Tel: 0277 210127

Pat Woodroffe, 9A Guildown Road, Guildford, Surrey, GU2 5EW. Tel: 0483 63299