

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.

Resurrection

The Bulletin of the Computer Conservation Society

ISSN 0958 - 7403

Number 5

Spring 1993

Contents

Editorial

Nicholas Enticknap, Editor 2

Guest Opinion

Maurice Wilkes 3

Society News

Tony Sale, Secretary 4

The Williams Tube Revisited

Tony Sale 5

Andrew Booth's Computers at Birkbeck College

Andrew Collin 11

The Origins of Packet Switching

Derek Barber 17

Altair and After: the Original PC Revolution

Robin Shirley 23

Letters to the Editor

32

Working Party Reports

34

Forthcoming events

40

Editorial

Nicholas Enticknap, Editor

Welcome to the latest issue of *Resurrection*, which we are pleased to say has taken a much shorter time to produce than its predecessor did. We have taken various steps to speed up the production process, and these are now bearing fruit. The next issue is already in preparation as I write.

In the meantime, there is much catching up to do. This issue is the largest to date, and includes edited versions of four talks which took place in the winter and spring of 1991-92. They are varied both chronologically and technically, spanning computing developments from the Williams tube of the forties to the pioneer personal computers of the seventies, and encompassing some interesting fifties developments as well as our first telecommunications feature, on the evolution of packet-switching.

In this issue we have included a couple of letters from readers who were spurred into epistolary activity by the articles in our last issue. Letters are always welcome, especially from readers who are rarely able to attend our Science Museum gatherings.

The Society has been very active since the last issue, and details of some of these activities are to be found in the Secretary's Society News piece and the working party reports. The most notable new developments are the formation of a sixth working party to restore the Science Museum's unique Elliott 401, and the creation of a so far unnamed branch of the Society in the Manchester area.

Work has also been proceeding on the plan to acquire Bletchley Park with a view to establishing a museum of computing and cryptology on the site. The Bletchley Park Appeal, launched in July, has proved a great success, with good press and TV coverage.

The media publicity was influential in persuading British Telecom to withdraw all their planning applications for the demolition of the wartime buildings in the Park. Negotiations have now started between the Bletchley Park Trust (of which our Secretary is now a member) and both Property Holdings and BT about the purchase of the Park.

One activity not reported in this issue is the diligent archiving work run by Harold Gearing with the assistance of two members of the Science Museum staff, Susan Julian-Ottie and Helen Kingsley. We hope to provide a detailed account of this work in a future issue.

Guest Opinion

Maurice Wilkes

It gives me great pleasure to be invited to provide a guest editorial for *Resurrection*.

As the first President of the British Computer Society, I find it particularly pleasing that the BCS has, together with the Science Museum, formed the Computer Conservation Society to document and illustrate with working computers the history of the now recognised profession of computing.

The explosive expansion of computing into all walks of life has been quite breathtaking. Such a rapid expansion has left much of the historical aspects poorly recorded, and unfortunately many important artefacts have been destroyed through ignorance of their true historic worth.

The Computer Conservation Society offers a mechanism whereby this can be and is being remedied. It will however only succeed if the people who worked on hardware, software and systems design look out all their old documents and records, including program tapes, and place them in the safe keeping of the CCS for proper archiving and preservation. We, who designed, built, maintained and programmed early computers, owe it to ourselves and future generations to see that our endeavours are properly recorded and archived.

The Computer Conservation Society must also redouble its efforts to involve more young people in its activities. Here the CCS's pioneering work on simulators has an important part to play. Many young people relish the challenge of producing complex interactive graphics on a modern PC. That these graphics emulate an early computer is an added challenge. It provides a bridge over which young computer professionals can cross into an alien world of thermionic valves, high voltages, early transistors and esoteric memory systems.

The concept of working with old computers owes much to the vision of Doron Swade, the Curator of Computing at the Science Museum, and to the support received from Dr Cossons, Director of the Science Museum. The partnership between a museum and a professional body has worked extremely well and may offer a useful model for other professions. May the Computer Conservation Society continue its pioneering work for many years to come.

Society News

Tony Sale, Secretary

A major development has occurred since the last issue which has made the future of the Society look rather more secure.

The Science Museum has instigated its Computer Conservation and Archiving Project, which substantially strengthens the relationship between the Society and the Museum. This is the initiative of Suzanne Keene, who joined the Museum as Head of Collections Services in the middle of 1992.

Essentially what has happened is that the Museum is now using the activity of the Society as a model for the procedures to be followed with similar projects, whether involving computers or other artefacts. In return, the Museum is providing greater support to the Society, both financially and in the provision of resources.

This has helped, for instance, with the archiving project being managed by Harold Gearing. Harold is now being assisted by Susan Julian-Ottie, who joined the Museum staff at the beginning of November. This means that the archiving project can make quicker progress—essential as the amount of material in our archive is increasing daily as people generously donate us their own material.

The most obvious sign of the new relationship is the decision to restore the Elliott 401. This machine has been waiting patiently in the corner of our working area for a couple of years, but work has been unable to begin until recently. The first meeting of the new Elliott 401 Working Party was held on 22 September, under the chairmanship of Chris Burton.

The emphasis is very much on conservation in addition to restoration. In this respect the Working Party is being assisted by Helen Kingsley, another to have joined the Museum staff over the past six months.

The Society is expanding in other directions, too. We have been approached by the Manchester Museum of Science and Industry about a Pegasus that they own, and the outcome is that a north-western branch is being set up to organise a restoration project for that machine. It is, incidentally, believed to be the first Pegasus ever delivered.

We hope that this branch will also provide a focal point for members in the north of England who are usually unable to come to our Science Museum meetings and seminars.

The Williams Tube Revisited

Tony Sale

The Williams tube was invented by Freddie Williams in 1946. During the next two years Tom Kilburn refined the technique, and used it to produce the world's first stored program computer—the date was 21 June 1948. This article describes the development of this interesting memory technology which was used in all the world's most powerful computers during the period 1949-54.

Memory has been a critical consideration for computer designers since the beginning. Babbage, in his proposal for the Analytical Engine, was very ambitious, talking about 100 memories of 50 digits each—though I think the biggest design that he actually produced was 16 x 50. He had certainly appreciated the need for memory, though, as did Vannevar Bush. He produced the designs in about 1936 for what was an exact equivalent of a modern VDU; and he put into that a very large memory. That's an important point in the development of ideas for what you needed in order to make a computer.

ENIAC had 20 memory locations, I think of 20 digits each—all valve memory. The original EDSAC proposal by Von Neumann and others was for 4000 words of 40 bits; Turing's ACE proposal was for 256 x 30 bits and Zuse produced a mechanical memory of 1024 x 16 bits.

So all those pioneers had appreciated the importance of memory, and there were lots of attempts to produce different sorts of memory. I've classified their broad characteristics.

The first is quantisation—that is, the number of discrete levels that are stored in any one location. Babbage was working on a decimal notation and Zuse was working on a binary notation. Binary is better as the distance between adjacent points can be kept further apart.

You can then have a sort of memory which is position dependent—this is what the electrostatic memories are. Position dependent can either mean in time or space. An example of a time-dependent memory is an in-flight memory where you are storing things by the fact that it takes a long time to go round a cycle, so that storage occurs during the time that it's in flight in the medium. A space-dependent memory could be an individual spot on a screen or an individual ferrite core.

Then you've got the problem of read-out — whether it's destructive or non-destructive. Babbage's read-out on his wheels was a destructive read-out (he had to restore them), but the Scheutz engine had a non-destructive read-out. The downside of that is that you had to differentiate between one of 10 quite close levels in order to find out which one it was actually set at, and that made it slightly unreliable.

The roots of electrostatic memory are in radar and the need for moving target identification. The problem with radar, as it got more and more powerful and got more and more echoes, was that you got more and more clutter. The difficulty was to sort out the moving target from the fixed background clutter. Quite early on people said: well why don't you just store the background clutter because it's not changing very much, and cancel it out on the next scan?

So there was an impetus to have a storage system which would hold all the information on a radar trace for one scan period and cancel out the fixed echoes. You should then get the moving one showing up.

One of the first ways tried was to use iconoscopes, because TV was well developed by the outbreak of the war and iconoscopes were easily available. The idea was to write a pattern and then compare the two patterns by reading them back again. That actually worked for MIT but when it was tried for storage for digital computers there was a problem.

The basic problem with the iconoscope or any TV artefact was that you could tolerate a high amount of noise because the eye does not notice it due to integration over time. So what looks like a nice static picture can actually be very noisy and have dropouts and bits missing.

So although iconoscopes gave perfectly good pictures for studio presentation of video information, when you get down to saying 'what's on that little bit there?' the answer was that in fact it fluctuated an enormous amount, and depended on all sorts of factors which they couldn't control accurately. So although iconoscopes were OK for doing a broad brush thing like TV or moving target identification in radar, they were never successful as a storage system for information for a computer.

Freddie Williams was working on radar at TRE, and he went to the States both in 1945 and 1946 to see what they were doing at MIT on moving target identification. He saw all the work they were doing and he came back to Malvern and set up a two tube system. But while he was in the States he also worked on the Waveforms book.

In that book is a description of the moving target cancellation using a

cathode ray tube. By this time it had been realised that you could store charges on the inside phosphor face of a CRT and by suitably modulating the velocity of the beam you could read out information.

When Williams came back to Malvern he then started setting up the two tube experiment, which involved reading from one tube, storing on another tube, then reading back from the second, so cycling round on two tubes. While he was doing this he discovered the anticipation effect which I'll explain in a moment. That led to continuous regeneration and to his first patents in December 1946.

Williams wasn't the only person working on electrostatic storage. It was realised that electrostatic storage had the potential for very fast storage as you could switch an electron beam very quickly from one point to another on a screen. It had the potential of random access anywhere on a screen to pick up what had been stored there previously. That was the reason it was felt to be so important; there were no other technologies which at that time could match it for potential speed.

So how did the Williams tube actually work? What Freddie Williams discovered was the anticipation effect; this is an effect which relies on secondary emission from the inside phosphor face of the CRT. It was found that for accelerator voltages of 500 to 2000, there was a peak at 1500 volts where the amount of secondary emission exceeded the primary current of the beam. There was an amplification effect, which was eventually used in photomultipliers and other things.

But the whole surface of the phosphor stabilises at about -20 volts with respect to the final anode (the aquadag round the CRT). When a beam strikes the CRT it charges it $+3$ volts or thereabouts with respect to the mean level, so that sits at -17 volts. The secondary electrons are all part of the stabilising cloud, which stabilises the voltage on the inside of the CRT. Williams found that making a spot move across the screen produced this 3 volt charge but as it moved to the next point the secondary emission from the next point filled in that charge again, so it neutralised it.

So the net effect is, if it's moving at a certain range of velocities, it will be stabilised, but at the end the beam is switched off. At that point there is a blank at the next part of the sawtooth: the last hole dug is never filled in again for there is no further spot. Therefore that leaves a charge at the end of the trench which is filled in by the secondary emission.

If you rescan it with another beam (which is on all the time in a plate on the front of the CRT) you get a waveform. This induced charge is caused

by the incident beam striking the +3 volts hill rather than the plateau of the rest; and that induces the charge in the plate on the outside. Because that occurs before the beam is switched off, you can actually use that pulse to switch off the beam which is now travelling which was an inspection beam.

So now you can regenerate because now you can actually anticipate that it's going to change there: this time round you switch it off and it regenerates. It's called the anticipation effect because the presence of that signal anticipates that the beam is about to be switched off.

There were various schemes tried; the original one was the dash dot system, but the difficulty was that that was rather sensitive to flaws and defects in the face of the CRT. It might work in the lab but once you go into production and talk about 2048 bits on the face of the tube, every one of those has got to work reliably.

Tom Kilburn found that there were problems with the manufacture of CRTs. In particular what happens is the aquadag coating is squirted on the inside of the tube neck to make the final anode, and very often some of that, in the later processing of the phosphor, would detach and sprinkle over the phosphor. So they had cases where you could move the matrix just slightly and it wouldn't work, then move it again and it would. This was very critical for the larger number of bits that they were trying to store later on in using the tube.

One of the ideas—I'm not sure who produced it—was to go to the focus-defocus method. You wrote a defocussed large area which dug out the charge, and then you inspected it with a fine spot to see whether in fact it had been dug out or not. So you were looking at either a diffuse spot or a spot which had been previously put there. The difference between those two was much greater, and much less sensitive to individual small flaws in the screen. So I believe that was extensively used in the latter days of the memory system.

There was another version of that which was the dot-circle one. I think the dot-circle method was used more in the States than the UK.

Tom Kilburn and Freddie Williams got a lot of flack from people because there was no real theoretical background for the Williams tube. It just bloody well worked! This always annoyed people because they couldn't prove it didn't work.

In fact in, I think, 1954 one of Tom's PhD students actually did proper research on how it worked and what was behind it all. The experiments

he did then were with the double dot method. What that means is: you fire a dot and then if you want to store a 1 you displace from that original position and fire another dot, and the debris from the second one fills in the first one. That is a more controllable situation from the point of view of measuring exactly what is going on than the dot-dash system.

The Williams tube was used extensively. It was actually used in the fastest computers in the world in that period for about five years from 1949-54, before ferrite cores arrived. ILLIAC was the fastest computer in its day, and that had 40 Williams tubes; it was a parallel machine 40 bits wide, one tube per bit, 1K words.

The storage system was used in the Manchester prototype and the Mk 1, and also with parallel architecture in TREAC at Malvern. In the States the important systems were the IAS Princetown machine, the IBM 701 and ILLIAC.

The sort of cycle time that Tom Kilburn was using was usually around 10–12 microseconds. His was a serial architecture in that he read the bits serially across a tube, so you selected a word and then read the bits out of a word. I think the IBM one got down to about 5–6 microseconds, but it was typically twice that.

There were always tremendous arguments between the Manchester people, the Cambridge people and the NPL people as to which was the fastest machine. Although Tom Kilburn could get there quickly, he then had to read out at 12 microseconds a bit. The other people who were doing on-the-fly ones had to wait a long time for the information to come round the tube, although the clock speed of Pilot ACE was a lot faster at 1MHz.

Manchester received royalties from a large number of organisations through licensing of the technology, one of which was IBM. Because there were a lot of royalties, that started lawsuits.

There was a claim by Eckert which implied that he had invented the electrostatic storage system. To stop that the NRDC mounted an interference action in the American courts. I have obtained from the National Archive in Manchester a copy of the briefing paper to the British Council (representing NRDC) on it. It is fascinating reading as it's all there, laid out with all the affidavits and all the dates. I'm glad to say that NRDC won the case, and it was deemed to be Williams' invention.

Because of that IBM licensed the Williams technology and they put it into the 701. I'm not sure how many of the 701s had Williams tubes because there was a transition period when ferrite cores came in, and later

701s were shipped with ferrite cores, but certainly some of the 701s and 702s went out with Williams tube memories in.

That led to some interesting research by IBM into the engineering of the Williams tube storage system. One of the things that people were worried about — particularly the IBM people — was a thing called the read-around ratio.

Because the charge leaked away slowly within the face of the cathode ray tube and gradually stabilised back to its -20 volts, you had to refresh the data. The refresh time was only a few milliseconds, but the decay time was about half a second.

It also meant that if you didn't revisit a given site within a certain time (the read-around time), because you were reading other bits on the tube, the sort of background hash generated from that would fill in the ones that you hadn't visited. So the fact that you were reading information from a tube and writing to it at various places caused a degradation of places you hadn't visited.

So on the 701 they used an interleaved address read and write method so as to avoid going to adjacent places on reading and writing. This technique was also used on core stores to try and reduce the amount of crosstalk.

The papers I've read about the 701 indicate that they were very worried about this degradation of the information on the tube. They were also worried about the uniformity of the electrostatic effect across the face of the tube.

On the first CRTs they suffered from all sorts of imperfections, and they found it was pollen dust — they were out in the country somewhere and they were making these tubes in a factory where there was a high pollen environment. It was basically all quality control. It didn't matter on your radar scan, which was what they were mostly used for initially, or on TV tubes; but when you were worried about an individual bit standing on its own, you needed a different level of quality control.

This article is an edited version of the talk given by the author to the Society at the Science Museum on 23 January 1992. Tony Sale is Secretary of the Computer Conservation Society.

Andrew Booth's Computers at Birkbeck College

Andrew Collin

A D Booth is widely recognised as a pioneer of computing: he originated the Booth algorithm, and he was one of the first producers of computers in this country, or indeed the world. This talk does not attempt to add to what is already widely known about his career and scientific work, but is confined to personal reminiscences.

Half a lifetime has passed since 1957. To set the scene, imagine a world where World War 2 was a recent memory; Stalin had died only four years before; the Cold War was at its height; Britain had just been defeated in Egypt; it cost just 2½d (1p) to send a letter; steam trains still ran; there was no colour TV; and anyone with a qualification had no difficulty in getting a job.

In 1957 most people had never heard of computers; many of those who had believed the word referred to a calculating clerk, or perhaps a mechanical gadget to help you shoot down an aeroplane.

In the final year of my engineering degree at Oxford, I came across one of Andrew Booth's early books, 'Electronic Digital Calculators'. I was intrigued by his ideas and arranged to visit him at Birkbeck College, where I was taken on as a research student.

At this time Booth was head of the sub-department of Numerical Automation, a section of the Maths department of Birkbeck. The sub-department was housed in an old wartime emergency water tank, which you approached by rickety wooden steps leading down from the street behind the College.

The sub-department had a small population — Booth, his wife Kathleen who developed many early ideas on programming, a secretary, and a few research students working in such diverse fields as linguistics, character recognition, crystallography and the behaviour of thin films. Booth himself had a country house at Fenny Compton, where he did much of his work. He tended to be in College only about three days a week.

The sub-department was equipped with two computers of Booth's own design and manufacture — the APE(X)C and the machine I made most use of, the MAC 1. Both machines shared the same architecture.

The MAC 1 was built into a frame about the shape and size of an

upright piano. The base was full of heavy power supplies with large transformers. The 'keyboard' was a sloping panel with neon lamps to show the internal state of the machine and a row of buttons for data input. The logic was built on to vertical panels, with the valves out of sight facing backward and the other components to the front so as to be immediately accessible.

The machine weighed about 400 lbs and used about 230 thermionic valves, mostly double triodes.

From the programmer's point of view the MAC 1 had a simplicity which is only beginning to be approached by the latest risc architectures. The cpu had two 32-bit registers, called the accumulator and the register.

The main memory was a rotating magnetic drum with 32 tracks. Each track held 32 words of 32 bits each. Ten bits were enough to address any word in the memory.

Input was through a five hole electromechanical paper tape reader of Booth's own design, and output was through a standard five hole 10 cps paper tape punch, as used in those days for ticker tape.

The machine used 32-bit instructions, and had a two-address order code. The instructions were arranged starting with the two addresses, followed by a function field, a counter field, and a vector bit.

The first address generally gave the data address, and the second indicated where the next instruction was coming from. Each address was five bits of track number and five of word number on the track.

The function field, containing four bits, could specify the following instructions: load to accumulator from memory; add memory to accumulator; subtract memory from accumulator; AND and OR operations; store accumulator to memory; multiply; rotate right using accumulator and register; input to accumulator from paper tape; output from accumulator to paper tape punch; conditional branch on top bit of accumulator; stop.

The machine was serial. Each revolution of the drum took 32 'major' cycles, and inside each major cycle there were 32 minor cycles or bit pulses. A simple data operation such as a 32-bit addition took 32 minor cycles. However, the number of cycles actually used for any operation was controlled by the counter field, which specified a six-bit starting value. The operation would be halted as soon as the counter overflowed.

For most commands the 'correct' starting value for the counter was 32, but for shifts any value could be used sensibly.

The vector bit specified a vector operation, which meant that the com-

mand was repeated for the whole revolution of a drum, using each memory location in turn. This was chiefly useful for multiplication, where the command actually specified one stage of the Booth algorithm.

The complete multiplication was done by loading the multiplier into the register, writing the multiplicand into every location on a given track, and executing a ‘vector multiply’. The operation would take exactly one drum revolution and leave a 64-bit product in the arithmetic unit.

As far as I am aware, the APE(X)C and MAC 1 machines were the first to incorporate the Booth multiplication algorithm, which gives correct results when multiplying twos complement signed numbers.

Much of Booth’s early work was in crystallography, and involved a great deal of calculation with desk calculators. Multiplication was done by adding the multiplicand repeatedly for each digit of the multiplier, and shifting the partial product one place after each sequence of additions.

A well known trick in those days was to speed up multiplication by using subtraction as well as addition. For example, a string of nines in the multiplier could be handled by subtracting once, shifting along several times, and adding.

Booth formalised this observation and applied it to binary multiplication, where it led to a remarkably simple rule:

- Examine each pair of digits in the multiplier, creating the first pair by appending a dummy ‘0’ at the least significant end: then
 - if the pair is 01, add the multiplicand;
 - if the pair is 10, subtract the multiplicand;
 - otherwise, do nothing.
- Shift both partial product and multiplier one place right, allowing the next pair of digits to be examined.
- Repeat as many times as there are digits in the multiplier.

The machines did not have a division instruction, although Booth was reserving a function code for this purpose. Division was done by program, and took a substantial fraction of a second.

Since the memory did not have random access characteristics there were obvious advantages in being able to place successive instructions in the ‘best’ place on the drum. This implied that each instruction needed to carry the address of its predecessor.

With the two address format there was no need for an unconditional jump. The conditional jump tested the top bit in the accumulator, and used the data address instead of the successor if this bit was set.

Testing for zero was clumsy. You had to test for zero-or-positive, then subtract one and test again. Then you had to add the one back to recover the original value.

Using arrays was awkward without an index register. To access successive elements, you had to increment the address in the instruction itself.

Programming tools were primitive. The main aid to programmers was the coding sheet and the pencil. To code effectively you had to know how long each instruction would take, and place each instruction at the best possible place: this was called optimum programming.

When the program was written it was coded (that is, each group of five bits was translated into the corresponding baudot character) and punched on a device originally intended for telegrams. Then the tape was read into the computer using a loader which lived on track 0. More often, however, the program was keyed into the computer by hand because the tape reader was not working reliably.

Another aid to programming was the library: we had a number input routine which lived on track 24, decimal output on track 26, and so on. The machine was innocent of any symbolic assemblers, but you could single-shot your program to find out where it was going wrong.

In terms of hardware, the machine must have been one of the most primitive ever constructed. However the technology available for computing was equally antediluvian.

The transistor had only recently been invented, and was not reliable at high frequencies. No transistors were used in MAC 1.

Logic gates were made from germanium diodes. These components are notoriously heat-sensitive, and on warm days we had to aim an electric fan at the innards of the computer to make it work at all.

The registers in the cpu were Eccles-Jordan flip-flops which used one double-triode valve per stage. A cascade of these stages was able to shift a binary pattern without any intermediate storage. Each flip-flop was linked to the next one by a couple of small capacitors, and when all the stages were forcibly switched to zero by a huge, very short 'shift' pulse, each stage was then able to take on the state of its predecessor.

I never understood how this worked, although the process has variously been described as 'magic' and 'stray capacitance'. In any event the system

was extremely sensitive to the size and duration of the shift pulse, which had to be adjusted, it seemed, even to suit the weather.

Nowadays all shift registers are designed with two phases, so that intermediate storage of the data is assured, but in those days valves cost the equivalent of £20 each and were extremely unreliable, so the best design was the one which used the smallest number of valves!

Over prolonged use, valves degrade in variable and unpredictable ways. Quite often a valve which failed in one circuit would work in another. This led to a popular method of fixing faults: take out all the valves and plug them back in at random!

An interesting feature of the machine was its drum store. It consisted of aluminium, coated with a layer of magnetic oxide. Flying heads had not been invented, and each of the 32 heads was attached to a solid brass pillar. Its exact spacing from the surface of the drum had to be adjusted by turning a screw. Since the correct spacing was only a few microns, this was a delicate operation.

We would set up the computer to write and read alternately, and examine the output of the head on a 'scope. At first there would be nothing: then, as the head approached the surface, information would begin to come back. If the screw was turned a tiny fraction too far the head would crash with a loud clang, and strip the oxide from the surface of the drum. That track would have to be abandoned until the drum was next recoated.

The drum also provided timing signals for the whole machine. Two special tracks had notches machined into them, and the signals from the heads generated the major and minor cycles, respectively.

Track selection on the drum was done by a relay tree. They were the fastest relays you could buy at the time, and on a good day they could switch tracks in only six major cycles. The delay had to be taken account of in programming: if you tried to read a word from a different track with a 'gap' of less than six words, you would get all ones.

Sometimes a relay stuck, and some tracks were inaccessible. If necessary we would reprogram the machine to avoid those tracks.

We were always interested in getting the machine to run faster. The overall speed was governed by the speed of the drum, which was driven by a synchronous AC motor at 50 revolutions per second. On one occasion I bought a second-hand motor-alternator set and installed it with a smaller pulley on the alternator, so it generated current at 75Hz. This speeded up the computer by 50%!

It is interesting to compare MAC 1 with present day technology. This talk was written on a 386SX PC. The table shows how primitive MAC 1 was in comparison.

	MAC 1	PC	ratio
RAM	8 bytes	2 Mbyte	250,000:1
Rotating memory	4096 bytes	40 Mbytes	10,000:1
Speed	50 ips	12 mips	250,000:1
Mean time between faults	30 minutes	5000 hours	10,000:1
Mean instructions between faults	100,000	10^{14}	$10^9:1$

Looking back, what did we achieve? In terms of practical computation, very little. All attempts at serious work were frustrated by the frequent failures of the machine, and the lack of credence one could attach to its results.

As a research tool, however, Booth's machines were highly successful. They were used to develop and demonstrate numerous programming techniques, and to educate a whole generation of research students.

One feature of the work at Birkbeck puzzled me for years. At the time we were doing our best to use MAC 1, effective and reliable computers such as the Ferranti Pegasus were already available. Why did we bother?

On reflection, it now seems clear that Booth's machines were early ancestors of the mini or PC. Pegasus was the supercomputer of its day. One machine served a large organisation, was tended by a full time professional staff of operators and programmers, and cost about £50,000 to install (say £1 million in today's money). But the total component cost of the MAC 1 was only a few hundred pounds, and the machine could be assembled by a skilled technician in about six weeks.

Booth's aim was to build a machine which could be afforded by small companies, colleges and even schools. If the machine had been more reliable it would have allowed a working knowledge of computers to be spread earlier and more widely than in fact occurred.

This article is an edited version of the talk given by the author to the Society at the Science Museum on 28 November 1991.

The Origins of Packet Switching

Derek Barber

I remember Donald Davies at NPL coming into my office one morning to discuss the future research programme. I think that would have been probably late 1965. His feeling, he said, was that data ought to be handled by a network rather in the way that parcels or packets are handled in the postal system.

At least that's the way I remember it. Donald tells me his recollection was that the packet idea came a bit later, but I think he's wrong. I'm sure he did mention packet at that first meeting. At the time he initiated a small survey as to how the word *packet* would translate into most of the world languages in order to judge its suitability. The result was generally favourable except for Russia where it was already in use as a data block in a link. But he decided that that wasn't much of a constraint so *packet* it was.

But let us start at the beginning. I started work in the Post Office Engineering Department. I did the rounds in the workshop with different lads and so on — jolly good ground work. I actually worked on the London to Birmingham one inch cable for television and I remember the night when the Birmingham transmitter opened in November 1949. The line went down five minutes before the show was due to go live. At least we thought so — eventually it turned out the film had broken in the telecine so they had to rewind it before they could start. The network was perfectly all right.

I eventually got a degree (part 1) in evening classes, then got through the open experimental competition. I got put into research branch and worked in the war room under the ground testing contacts for nearly a year. Then I got put into RC4-1 which was a marvellous division, working on pulse and bar testing of television links. Then I got my full degree, applied for the position of open scientific officer, and got sent down to NPL.

The Director was Bullard and at that time when I went for the interview about March 1954 there were two sections: Control Mechanisms in Metrology, and Electronics in Maths Division. These were run by Dick Tizard and Frank Colebrook.

At the interview I sat down in the large conference room at the big oval table, and they asked where I had been working. I said Dollis Hill

and I had been working on vestigial sideband transmission. Bullard said “Just a minute, that’s a technical term. You must explain it.” Colebrook leant forward and said “Perhaps I can handle this one, Director.” And he chatted for about 10 minutes about modulation methods and the like. Then Bullard said “Christ, I’m due up in town in about 10 minutes. I’m late already. Nobody else has got any questions have they?”. And that was it!

NPL formed a division with Dick Tizard as Superintendent, and I got working on guided weapons data processing. Tizard then went to LSE and Ted Newman became acting Superintendent. I remember Ted ringing me up one day in 1958 and saying “I’d like you to represent us on the IEE Measurement and Control Section.” I said “Oh, I don’t think I can do that.” “Oh yes” he said “I’m confident you can.” So in the end I did it. Yet it was some years afterwards I realised he was just trying to find a mug to take on the job.

We made a digital plotting table. I built a very high speed analogue amplifier-multiplier. Then we started work on an alcohol still as part of our work on adaptive control under Percy Hammond. I got my PSO promotion then. That was in 1963.

I meant to say something about that plotting table. I got some transistors in from BTH to build the plotting table. It was an interesting design because it had a binary point and about 10 places of binary fractions. I don’t think we ever had a situation where all of the transistors were working at one time. That created a problem of accuracy as you can imagine.

I went to the States in 1963 and spent four weeks there, which opened my eyes. I went to MIT and saw Project MAC; I saw the PDP-1, and the Sketchpad work that Ivan Sutherland was doing.

As a result of the distillation column instrumentation we got working on a data processing system with standard interfaces and data transmission, and that paved the way for the network.

In 1965 Dunworth was Acting Director, Uttley was still running Autonomics Division, and the SPSOs were Percy Hammond, Ted Newman and Donald Davies. I was under Percy doing adaptive control and the NPL Standard Interface.

I became chairman of BSI DPE 10, which is how I found out about politics in standards. Eventually at a Berlin meeting in 1969 the BS 4421 British Standard Interface was being offered as an international standard.

There was a dead heat and the chairman gave his casting vote against it.

By October 1970 I gave up the BSI work because other things were beginning to build up like the Data Communication Network. But this 'packets' meeting with Donald Davies in mid to late 1965 came about partly because of this work.

Donald himself went to IFIP 65 in the States and he came back from that with the view that you ought to handle data in the same way that you handle data in a time-share machine, with time slicing, source allocation and so on. I think that was basically where the thoughts came from. On the 18th March Donald gave a lecture at NPL but there are one or two papers before that which have never been published.

So Donald was going public then on packet ideas, and we formed Autonomics project No 6 on Data Communication early in 1967. So that's the background.

Soon after that Uttley left and Donald became superintendent (I think). Percy was one of the SPSOs and Ted was the other one. I got the SPSO in 1969 and I picked up the other half of the division. For a while I thought we had a marvellous team, the three of us used to meet on Monday mornings in Donald's office and really kick ideas around. I look back upon that time as a time when things were just right.

Out of that came the data communications work. Roger Scantlebury took over that. Keith Bartlett was with the hardware team, and for software we had Peter Wilkinson, John Laws and Carol Walsh. Pat Woodroffe stayed with Alan Davies for a while on the BS 4421 and was instrumental in the big display. I think Brian Wichman wrote the cross-compiler for the KDF9. Anyway we had a Mk 1 and a Mk 2 system. The Mk 2 was an altogether better design.

We had a Modular 1 to run Scrapbook. Maths Division had the KDF9. There was a PDP-11 front end on which ran the Edit service which was eventually made available on the network.

At Donald's talk in March there were over 100 people including 18 from the Post Office. I mention them because they featured quite a lot in what happened afterwards as far as Roger and I were concerned. There had been a paper written by someone from the Rand Corporation which, in a sense, foreshadowed packet switching in a way for speech networks and voice networks but nobody knew anything about it and certainly it didn't enter into our thinking at all. Eventually Donald wrote an internal paper which was really his lecture polished up.

Then the Real Time Club was formed early in 1967. Rex Merrick and Stan Gill were very important then, and Donald of course. In the meantime Roger went to the ACM Symposium on operating systems at Gothenburg. There wasn't a conference about networking; of course the subject hardly existed, so the operating system seemed to be — since it was timesharing — the right place to go.

Anyway Roger went and gave a paper called “Digital Communication Networks for Computers giving a Rapid Response at Remote Terminals.” Larry Roberts had extended the concept of a support graphics processor to the idea of a network, and he was then talking about multiple computer networks and inter-computer communication. Roger actually convinced Larry that what he was talking about was all wrong and that the way that NPL were proposing to do it was right. I've got some notes that say that first Larry was sceptical but several of the others there sided with Roger and eventually Larry was overwhelmed by the numbers. That actually gave birth to Arpanet because Larry joined soon after and became responsible for it.

Events happen and it's difficult to get them chronologically right. But certainly by July 1968 the Real Time Club had organised this great big event in the Royal Festival Hall. We worked jolly hard to produce various bits of kit and so on and we were able to put on this show (either simulated or using real networks). This provoked a debate between Stan Gill and Bill Merriman, because the Post Office at the top level were all telephone people and it must have been very hard to take on board something very new.

In autumn 1969 there was the Mintech Network proposal. Roger and Donald came round to my house and spent an evening discussing this and putting ideas together.

In November 1969 I went to the USA and I saw the first three Arpanet nodes that were ready to be shipped out to the West Coast. I remember going round with a set of slides and giving a talk. When I came back the Director was a bit worried in case I'd been giving away all our ideas. Donald's response was that we had to tell people about it to get anything to happen.

Two and a half hours of KDF9 run time accounted for half a second of run time on the network — an amazing ratio really. It just shows how slow the KDF9 was.

Then came the Isorhythmic: Donald was doing work on controlling

congestion there. Then Costas Solomonides came to join us and did a lot of work on hierarchical networks in relation with Logica.

Then there was the Mark 2 NPL network software. The first lot had been written in assembler or something. Peter Wilkinson worked for five months and nothing really appeared except strange transition diagrams. Eventually the software got written and they loaded it all up and there were two bugs that they cleared in a day; and from then on it just ran.

Eventually it was all rewritten in this PL516. In doing that we had Ian Dewis who came and joined us and later went to British Steel where he got involved in their network. Then Alan Gardner was seconded from the Post Office to us, so we began to get people coming from outside. The reputation had got around of this interesting work going on.

There were also the attached services. Scrapbook I've mentioned, and Edit, involving Tony Hillman and Roger Schofield. Then we had a File Store built by people from CAP. Then we had a Gateway to EPSS (the timing is a bit uncertain). Moving on towards EIN now, John Laws was responsible for the EIN management centre; he's at RSRE now.

Just a few words about international things that were going on. I've already mentioned Arpanet and Larry Roberts. Telenet was basically a company to exploit Arpanet and Larry Roberts was its president. Eventually that was taken over by GTN. That was a recognised carrier so Telenet became respectable, and therefore a recognised public operator.

I mention Peter Kirstein at UCL with the gateway to Arpanet because over a period of time that's been quite significant. Certainly we used to use the Arpanet message service — electronic mail — through that gateway quite often. I did even when I was in the Alvey Directorate.

In Canada Bill Morgan built Datapac about this time. By now I was running EIN and I got involved a lot because of that. I also got involved in CCITT, SET and a whole manner of things.

I actually set up the first meeting between John Wedlake of the British Post Office and Rene Dupre of the French PTT which led to X25. There was a problem about virtual calls in EIN, so I called this meeting and that actually did in the end lead to X25.

A philosophic point about networks — if you make it properly dynamic resource sharing and it all runs much faster than any user, there's a high probability the user gets what he wants. But the PTTs are not happy about that, partly because of the background (the fact that the telephone network provides a connection and so on), and partly because they've got

an obligation to provide a guaranteed service of some kind. If they give you a circuit and you succeed in dialling it up then you expect to get 4 KHz; you don't expect $2\frac{1}{2}$. So their philosophy was very much that you had to simulate an end to end connection.

Rene Dupre was a man who believed in allocation of resources; so in the French RCP network he had buffers allocated at every switch per call. Arpanet had buffers allocated at the ends but in the French network they didn't even allocate it at the ends; they basically did it in what amounts to host computers.

But Dupre had got a guilt fixation about this and there were these meetings because the PTTs had to get agreement on X25. There's an interesting difference between PTTs and computer people when it comes to standards: the PTTs sell services, and you can't sell services if you don't have standards, which is why the bottom three layers of the ISO pattern and the top layers are hardly released yet.

But Dupre went around telling everybody that if you don't build it our way we won't get an agreement. And if we don't get an agreement, none of us will be in business, because we won't be able to sell data networks: that's roughly how it went.

The Cyclades network led to Transpac so the French PTT in the end got off the ground with a French network. Cezar who did that was involved in EIN with Logica.

The Spanish, dark horses, were the first people to have a public network. They'd got a bank network which they craftily turned into a public network overnight, and beat everybody to the post.

This article is an edited version of the talk given by the author to the Society at the Science Museum on 30 April 1992.

Altair and After: the Original PC Revolution

Robin Shirley

The almost accidental commissioning of the Altair 8800 micro-computer kit in 1975 to accompany a magazine series proved to be the catalyst that launched a snowballing personal computer movement based on machines that adopted its 100-way bus as a de facto standard.

This account chronicles the movement's explosive growth and examines the factors that fuelled it, the people and companies that were involved, and the populist, libertarian political and social ethos that it sought to promote.

Origins

What provided the impetus for the personal computer movement? Not the established computer industry, at least not directly.

There was a growing substratum of young, smart programmers and users, mainly in universities and colleges but also from industry and business, who had for years nursed a love-hate relationship with the crude, inflexible mainframe computers and hostile, autistic systems software on which their jobs ran, or, just as often, failed to run.

Of course there were also smaller and neater computers, but only the fortunate few got their hands on them and could feel they were masters of their own fate. These would usually be small, dedicated minis like PDP-8s, PDP-11s or maybe Novas, in a science or engineering lab.

So the primary issue was freedom from interference and frustration. The other main one was power.

The mainframe computer was an obvious symbol and concentrator of corporate and official power. It tended to reinforce all the tendencies to centralism that don't need any encouragement in any sizeable organisation. On the other hand, it could clearly also be a tremendously powerful tool if exploited effectively.

In those days a lot of mileage was got out of 'Grosch's Law', a rule you don't hear much of now, which proposed that the power of a computer system increased as the square of its cost. So computers seemed to mean more and more power for the big battalions, who arguably already had as

much power as was good for them.

Reactions to this situation varied. The general public tended to take a broadly luddite position — computers are essentially anti-people and should be curbed — a view that still has its adherents today.

Others felt in their bones that there had to be a better way ...

The story of the rise of the microprocessor and large-scale integrated circuitry is a familiar one, so I won't dwell on it. In essence, in 1971 a small team of mostly ex-Fairchild people led by Ted Hoff produced the first commercial microprocessor, a 4-bit unit — the Intel 4004 — meant to provide the driving logic for a Datapoint intelligent terminal, but pursued for other uses when the original order fell through.

The 4004 was doubled up to give the 8008, a primitive 8-bit processor which lacked, for example, any direct memory addressing instructions, and this was improved and refined in 1973 to give the Intel 8080.

The 8080A, an NMOS chip clocked at 2MHz, was the processor used in the Altair and its first generation successors and gave them their characteristic architecture. It had an 8-bit data bus, a 16-bit address bus (and hence 64Kbyte direct-addressing range), a separate 8-bit I/O address space that defined 256 'ports', a mixture of 8- and 16-bit registers, and a reasonably adequate order code of 78 instructions. Intel also provided a family of support chips that made it relatively easy to produce a complete 8080-based system.

Z80

It was followed up a couple of years later by an enhanced and extended version, the 8085, but by then most of the S-100 microcomputer interest had shifted to the Zilog Z80, produced in 1976 by a small group of engineers who had, after the fashion of Silicon Valley, split off from Intel to go their own way.

The Z80 ran at 2.5MHz and had an extended order code of 158 instructions (400 or so if you chose to count the different bit-manipulation orders separately), a single +5v power rail and a single-phase clock, so it was an altogether more elegant device than the 8080.

Best of all, its order code was an almost exact superset of the 8080's, so that a Z80 could execute nearly all 8080 programs unchanged. Its faster 4MHz Z80A version offered real power, and featured in most of the second generation S-100 systems. Eventually still faster versions appeared, the

6MHz Z80B and 8MHz Z80H, but by then other architectures had moved to the fore.

Meanwhile, another line of development and spin-offs led to the Motorola 6800, used in the South West Technical Products machines and in the Altair 680, and the MOS Technology 6502 (1MHz and 2MHz) that powered other important (non-S100) machines like the Apple II, Commodore Pet and BBC micro.

Altair

In late 1974, a series of computer construction articles on an 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 so the Altair 8800 was born.

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 called the Altair bus.

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

Meanwhile, improved clones started to appear, so that by August 1976 *Doctor Dobbs' Journal of Computer Callisthenics and Orthodontics* (DDJ) was calling it 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. In due course, some five years later, a cleaned up 8/16-bit version became officially standardised as the IEEE 696 bus.

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

Floppy disc systems were appearing too—at first rather bulky ones based on single-sided single-density (SSSD) 8-inch drives, holding a nominal 250Kb (kilobytes) per diskette. Soon, however, these were supplemented and eventually supplanted by a new 5.25-inch Shugart SA400 mini-floppy format. In SSSD form these stored about 175Kb. For several years the 5.25-inch drives continued to be scorned by 8-inch disc users in the kind of shallow partisanship that often afflicts technical enthusiasts.

Among the vendors of add-on disc systems was a Californian company called North Star, whose blue-painted drive cabinets became a common sight. Their popularity was based on reliability, low price and (especially) the fact that a somewhat spartan but efficient operating system (North Star DOS), accompanied by an excellent Basic (using BCD arithmetic and hence good at avoiding roundoff in financial calculations), was bundled free with their disc controllers. Floppy-disc North Star DOS was stripped for speed, with few concessions for convenience, and easily out-performed CP/M and indeed many hard disc systems too.

They 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 like AST and Everex, North Star was soon to use this experience as a springboard into producing complete systems.

Second generation

The period 1976-77 saw the arrival of a host of high quality second-generation designs. This was the golden age of S-100 systems, in which appeared classics like the Cromemco Z2, North Star Horizon, Vector Graphics MZ and Ithaca Audio (later Ithaca InterSystems) DPS-1.

The design of the classic small business or scientific microcomputer system crystallised as a 4MHz Z80A-based S-100 machine in a 19-inch cabinet with twin 5.25-inch floppy drives, running under CP/M.

Horizons in particular were very widely used in the UK, though outsold by Vector Graphics and (more marginally) Cromemco in the USA. The Horizon seems to be remembered with affection by all who used it, as an elegant, rugged and extremely stable and long-lived design. It has become part of the industry folklore that the engineer's console on the original Cray 1 supercomputer was in fact a rack-mounted Horizon.

The Horizon motherboard design, with its input/output circuitry mounted on a rearward extension of the PCB, was notable for its far-

sightedness and provision for every variant of asynchronous and synchronous I/O or interrupt servicing that might be needed, which was specially useful for OEM applications where (like other S-100 machines) it could easily be built into a standard 19-inch equipment rack. This sort of use was quite common and accounts for many of the machines still active today.

Looking back over 10 years or so of servicing Horizons, I'm still impressed with how few design faults they had—I can really only think of two, both relatively minor and arising from an apparent blind spot on the part of its designer, who tended to disregard the long-term consequences of what happened to waste heat once he'd dispatched it to a heat sink.

CP/M

Just as significant to the success of the S-100 microcomputer as their hardware standardisation was the standard software environment offered by their CP/M operating system.

In 1973, Gary Kildall, a young software consultant at Intel, was fed up with trying to develop the PL/M programming language for microprocessor development systems on paper tape using an ASR 33 teletype, and so begged an ex-10,000-hour-test floppy drive with worn out bearings from the marketing manager at Shugart Associates, a few miles up the road.

However his attempts at interfacing proved abortive, and it was not until late 1974 that a colleague, John Torode, took an interest in his problem and completed a wire-wrap controller to interface the drive to Gary's Inteltec-8 development system.

Meanwhile Gary had put together a primitive disc operating system for the drive, and in due course (according to Gary) the paper tape was loaded and, to their amazement, the drive went through its initialisation and printed out the system prompt on the first try (legend doesn't record whether it also did so on the second try).

Gary named the operating system CP/M, which in early accounts stood for Console Processor and Monitor, but later became dignified as Control Program/Microprocessors. It was ported to two other (non S-100) microcomputer systems during 1975, and Gary continued to work on it in his spare time, producing an editor, assembler and debugger—ED, ASM and DDT (the style and nomenclature of CP/M were heavily influenced by DEC operating systems).

In 1976, IMSAI shipped a number of floppy disc systems with the

promise that an operating system would follow, but as yet none existed! Glenn Ewing, who was then consulting for IMSAI, approached Gary Kildall to see if he would adapt CP/M to fill the bill. Gary agreed, but so as not to have to change CP/M again as a whole to fit another computer system, he separated out the hardware-dependent parts into a sub-module called the BIOS (Basic Input/Output System), so that any competent programmer could then do the job.

This version, CP/M 1.3, was distributed by IMSAI with modifications as IMDOS, and Gary was also persuaded by Jim Warren, the editor of DDJ, to put it on the open market at \$70 a copy. Gary did this rather against his better judgement, since contrary to Jim Warren's assurances he was pretty sure that unlicensed copies would immediately become circulated.

However, to the amazement of sceptics, it was treated as a point of honour among the delighted S-100 users never to pass on their copies of CP/M, and the rip-off factor turned out to be practically nil. Gary Kildall formed a new company called Digital Research to support CP/M, and quickly became a millionaire.

CP/M was updated to versions 1.4 and 2.0, and then stabilised as CP/M 2.2 for several years. Late in the day, a version 3.0 for bank-switching systems with more than 64Kb of memory was produced, but by then 8-bit systems were on their way out, and, apart from a subterranean existence in the Amstrad PCW, it saw little action.

The standard software environment provided by CP/M proved to be the final catalyst that was needed. From then on, the main efforts of numerous independent microcomputer software writers could be directed into providing packages to run under CP/M, with the confidence that this would open up the entire market of S-100 systems for their products. The era of what was to become known as 'shrink-wrapped software' had arrived.

And what software! An explosion of very high quality programs followed, often written by some of the top big-machine programmers in their spare time, and sold at prices that were minute compared with their counterparts on minis and mainframes.

Whole new categories of software sprang up in this fertile environment. Interactive screen editors like Electric Pencil and Wordmaster led to MicroPro's WordStar, the first full microcomputer word processing program, which outperformed most dedicated word processors at a fraction of the

cost, and let microcomputers start recovering their investment by doing useful work from day one, something that seldom if ever happened with big machines.

WordStar sold a lot of microcomputers, and even more so did VisiCalc and SuperCalc, the first spreadsheet programs. Microcomputer programmers were now no longer just recapitulating minicomputer and mainframe software development, but breaking completely new ground.

The crude but promising database program Vulcan was transformed by Ashton-Tate into dBase II (so called to inspire confidence in its maturity — there never had been a dBase I).

IBM

In early 1980, elements within IBM started to sit up and take notice of the booming personal computer industry. The way it happened was that a year or two earlier, IBM had established a cluster of so-called Independent Business Units (IBUs in IBM-speak) with permission to act semi-independently, largely unfettered by the IBM bureaucracy, and a brief to break into new markets. Among these was the Entry Systems (Personal Computer) Unit, tucked safely away from the company's mainstream in Boca Raton, Florida.

In July 1980, Philip D Estridge, a divisional vice-president, was put in charge of a team of 12 and given a year to create a competitive personal computer. About 13 months later, in August/September 1981, shipments of the IBM PC started.

At this point it is helpful to remember that a number of obvious contenders among the large computer and electronics companies (Texas Instruments, DEC, Intel and Motorola, for example) had already made their move and failed dismally, generally through what they thought of at the time as doing things professionally, but what in hindsight looks more like typical big-company habits of inertia and over-pricing. DEC, for example, knocked the prospects of its Rainbow PC on the head in a classic bit of marketing over-reach by deliberately not providing a format program, so that users would have to buy all their floppy discs from DEC.

In consequence, it had become an article of faith in the PC movement that the big corporations stood as little chance of getting there as the dinosaurs had had of supplanting mammals. The bigger the less likely, it was assumed, so the longest odds of all would be against IBM.

So what went right at IBM? The big difference at Boca Raton was that the design team had a reasonably free hand and, crucially, that it included a number of computer hobbyists and hackers (in the original and proper sense)—people who already owned and were familiar with the existing personal computers—and that the project was allowed to adopt their open systems philosophy and reflect their user experience.

Except in one important respect (the 8088 with its 20-bit address space, which according to Peter Norton was urged on them by Bill Gates at Microsoft), the original specification of the IBM PC was intentionally a combination of the better (or at least well-tried) practices from the various 8-bit machines. It was perceived as such at the time, and this was seen by both users and reviewers as a virtue, since they were inclined to blow a raspberry at large computer companies who disdained to follow the custom and practice of existing users.

The base model had only 16Kb of RAM, plus 40Kb of ROM including an embedded variant of MBASIC, five expansion board slots and a cassette interface (!) for external storage. The RAM was only expandable to 64Kb on the motherboard, but to 256Kb using memory expansion boards - though not quite so far in practice, since 64Kb boards were the largest then available, and these had to compete with other boards for the five slots available, one of which would be pre-empted by a display adapter, and another for a floppy disc controller (and quite probably a third for a serial port).

The 5.25-inch floppy drives were at first SSSD and provided only 160Kb per disc, but IBM subsequently switched to double-sided double-density (DSDD) units holding 360Kb, mostly Tandon TM100-2's, as used, for example, in North Star Horizons. The standards of internal construction matched (but mostly didn't exceed) the best practice of S-100 manufacturers like North Star, Cromemco and Godbout CompuPro. This was enough to put it well ahead of most minicomputer standards.

The operating system PCDOS was a CP/M clone for 8086 processors bought in from Seattle Computer Products and hastily converted by Microsoft. Version 1.0 more or less worked, but was stark in the extreme. Moreover its command formats had been changed in the direction of Unix, just enough to make it tiresomely different from the quaint but familiar nomenclature of CP/M (which had in turn been inherited from DEC operating systems).

At this stage, Boca Raton was still undecided whether the PC would

primarily be a home computer or a business machine, so they hedged their bets. The two video cards offered—the CGA (colour/graphics monitor adapter) for games and the MDA (monochrome display adapter) for professional use—reflect this ambivalence. Also, as we’ve seen, a cassette interface was built in as standard, and the 320x200 colour graphics and 40 column text modes of the CGA adapter were chosen to match the meagre bandwidth of NTSC domestic TV sets.

As one might expect, the promotion and marketing of the IBM PC was impressive and the documentation superb. The Technical Reference Manual, which opened up the architecture to add-on board manufacturers, won special praise. Another essential break with IBM tradition, based on a study of Apple’s successful methods, was the use of independent dealers and franchised networks like Computerland, as well as direct corporate sales.

The price structure was also carefully pitched to undercut slightly its similarly configured rivals, at least in the US. Over here the time-honoured pound for dollar equation was not resisted, so that, like the Apple II before it, it didn’t get a look in as a home computer against the smaller but similarly configured 48Kb Spectrum at nearly a tenth the price.

In the event, the most important aspect turned out to be the way software vendors seized with both hands the opportunity to distribute on a standard disc format and, especially, to write for a standard (more accurately, two standards) of screen addressing—and the rest is history.

In little more than a year, sales of the IBM PC far exceeded those of all the S-100 manufacturers. The tone of the microcomputer world was increasingly set by the new cohorts of PC business users, and the personal computer revolution, as originally envisaged, could be said to have died of success.

From 1986 on, rocketing price-performance on the PC/AT-clone front quietly buried further development of S-100 systems.

This article is an edited and abridged version of the talk given by the author to the Society at the Science Museum on 27 February 1992. Robin Shirley is Chairman of the S-100 Bus Working Party.

Letters to the Editor

Dear Mr Enticknap,

May I please draw your attention to a small but quite significant error in Volume 1 Number 4 of *Resurrection*. It appears on page 21 (line 3 onwards) in the report of Tom Kilburn's talk to the Society about the Mark I computer. His memory is clearly at fault regarding the position and contribution of Sir Ben Lockspeiser.

Sir Ben was never associated with Ferranti Limited except as a customer. In 1948 he was Chief Scientist at the Ministry of Supply and very concerned that some of their development programmes were being held up by a lack of adequate computing facilities in the UK. With this in mind he visited Manchester University accompanied by Eric Grundy, the Director of Ferranti Limited responsible for Instrument Department which, as Kilburn said, was helping the University's computer project. Sir Ben wanted to assess the possibility that the computer project could be of practical assistance to him. He was favourably impressed and being a decisive and energetic man, promptly wrote to Eric Grundy, a letter intended to get things moving.

An ancient photocopy of a contemporary typed copy of that letter is enclosed herewith. Unfortunately a search carried out some years ago failed to unearth the original in the Ferranti files. However, that letter, dated 26th October 1948, effectively brought Ferranti Limited into the computer business and played a seminal part in the development of computers and computing in the UK.

The final sentence in Sir Ben's letter became an oft quoted phrase in the company.

Yours sincerely,

MH Johnson

Oxford

11 September 1992

Editor's note: the text of the copy letter referred to reads as follows:

Dear Mr Grundy,

I saw Mr Barton yesterday morning and told him of the arrangements I made with you at Manchester University. I have instructed him to get in touch with your firm and draft and issue a suitable contract to cover these

arrangements. You may take this letter as authority to proceed on the lines we discussed, namely, to construct an electronic calculating machine to the instructions of Professor FC Williams.

I am glad we were able to meet with Professor Williams as I believe that the making of electronic calculating machines will become a matter of great value and importance.

Please let me know if you meet with any difficulties.

Yours sincerely,

B Lockspeiser

Dear Mr Enticknap,

I have just read with great interest the article about the early days of Algol in *Resurrection*. It brought back many memories of my university days and early industrial career.

I began programming in 1968 in the sixth form at school, sending Algol tapes to the 903 at Medway Polytechnic. When I went to Leeds University in 1970 to read Mathematics and Computational Science, the teaching machine for first year undergraduates was also a 903. It was a 16K machine that had cost the university £22,000 in 1967!

From that machine we moved to a KDF9 and first experienced the joy of having filestore and on-line access! The latter was provided by the Eldon 2 operating system developed by Dave Holdsworth and others. I recall that there were two Algol systems (it was rumoured that two English Electric teams at Kidsgrove and Whetstone respectively had developed them, each in ignorance of the other). The Kidsgrove variety was a compiler, while Whetstone was interpretive.

After graduation I was absorbed into industry and converted to Cobol; but I was reunited briefly with Algol in the mid-1970s at the Jonas Woodhead Group who had a DECsystem-10 and were one of only two customers (the other being Whessoe in Darlington) who used it as a commercial and not a scientific machine. I recall computing Ackerman's Function in Algol to demonstrate recursion to a day release student.

Best wishes,

Yours sincerely,

Tony Peach

Telford, Shropshire

21 September 1992

Working Party Reports

Elliott 401

Chris Burton, Chairman

We formed the Elliott 401 working party last autumn, once the Science Museum's Computer Conservation and Archiving Project became an authorised funded project. Our objective is specifically to conserve and restore the Museum's 401.

This historic computer, an ancestor of Pegasus, is a unique one-off machine built by Elliott Bros in 1952 to prove packaged construction techniques and use of magnetostrictive storage in a complete computer system. It was demonstrated at the Physical Society Exhibition in April 1953 and subsequently was installed at Cambridge University where it was evaluated and modified by Christopher Strachey.

Later it was installed and used at the Rothamsted Agricultural Research Station until 1965, when it was donated to the Science Museum. So it has been in store for nearly 30 years.

Our work has two aspects—conservation and restoration. The former is concerned with the careful surveying of all the equipment, cleaning it, repairing damage to insulation, metalwork, paint finishes and so on, and generally bringing the machine to a stable and preservable state, as a normal museum artefact.

The restoration aspect will be concerned with making the machine work again functionally. The Working Party has been set up with three kinds of members—Museum staff with the knowledge and resources to do the conservation work, a small number of CCS volunteers with the experience and resources to tackle the restoration, and, importantly, the surviving original development team members, who act as consultants and sources of know-how.

The project is high profile from the Museum's point of view, and care is being taken to set a high standard of procedures in the spirit of the CCS aims to use voluntary expertise in conjunction with formal curatorial practice. A detailed plan and list of tasks is evolving, which will probably lead to an operational machine in about two years, depending on available resources.

We have held four formal Working Party meetings, which are the mechanism to get agreement on what to do and how to do it. These will probably

become monthly. In between we have had occasional days of preparatory investigation work.

But the main work has been the excellent conservation progress on the major units of the system. This is likely to take a further nine months. Already the base plinth and part of the top ducting has been completed, and the site at the end of the Old Canteen is beginning to take shape. Because we cannot use any of the units until they have been conserved, we have adopted a strategy of gradual re-commissioning using temporary sub-systems (particularly the power supply system) in order to get some restoration work going in parallel with conservation.

A significant number of original drawings exist in the Science Museum Library, which have been copied, but sadly some key documents such as the 'current' logic diagrams are missing. We will have to reconstruct such information by examining and recording the back-wiring.

We intend to attempt to rescue any information which may still be on tracks on the drum. We have not obtained any contemporary program tapes yet, but it will be some time before we will be in a position to consider running programs, though we are considering the desirability of a simulator.

So, a good start to an ambitious long-term project, thanks to the skill and enthusiasm of everyone involved.

Elliott 803

John Sinclair, Chairman

For some time now the processor has been suffering from an intermittent temperature-dependent store fault. It has taken a while to locate the problem, but the fault has now been found, and I am waiting for a replacement store read amplifier from the warehouse at Hayes.

The reliability of the film system has improved enormously over the past six months. It now works each time it is switched on, whereas previously new faults developed almost daily.

Readers may be interested in a statistical analysis of the faults we have encountered since the machine started running in October 1990.

The paper tape station has had 20 component failures, and we have also found two logic design faults.

The film system has had 36 component failures. We have also found one original wiring fault and three connections that had never been sol-

dered. (It is astonishing that these connections nonetheless worked perfectly throughout the machine's operational lifetime - it has taken 30 years to discover them!)

The central processor has had six component failures. Here also we found one connection that had never been soldered.

The high incidence of faults on the paper tape station and film system are due to the type of logic used, namely the Minilog potted logic element. They have proved to be much more unreliable than the logic elements used in the processor. This may be due to the fact that the transistors in the Minilog element are surrounded by a potting compound.

Our 803 emulator is now almost complete, as reported elsewhere. I have modified the software and hardware of a Z80 processor board, normally used to monitor telephone calls on a corporate switchboard, so that it now reads the 5-bit parallel character signals transmitted to the Creed teleprinter that is in the 803's paper tape station.

The Z80 converts the signals into a serial character data stream suitable for connection to the serial comms port of a PC. This facility allows the 803 to output data (normally copies of paper tapes) into a PC disc file for use by the emulator, and also provides an alternative means of paper tape duplication or archiving.

DEC

Adrian Johnstone, Chairman

We have organised our space within the old canteen, transferring some of our PDP-8 equipment into a store in the old School of Sculpture buildings which are nearby. This has made space for our PDP-11/20 and 11/34 systems which are being shipped up to the Museum from RAF Wroughton near Swindon.

The PDP-12 had has one tantrum, during which wisps of smoke appeared. The source of this fault has never been traced, but since it has not recurred we have decided to leave well alone.

The Open Day was a success, with the PDP-12 and a variety of PDP-8s entertaining our visitors.

Software and Emulators

Tony Sale, Chairman

Members of the working party have continued their work on developing emulators since the last issue. The activity is developing quite an impetus, though it is proceeding through the efforts of individuals rather than via formal meetings.

I have now acquired a 486-based personal computer, and am in the process of transferring the code from my previous Amiga, so as to facilitate further development of the animation mastering system.

Peter Onion's 803 emulator is now almost complete. We have received a good response from members who have sent us discs so that they can receive a copy. These discs will be sent out shortly.

Work has started on developing other emulators. One of the most interesting is a third year degree project being undertaken by Neil Mitchell of King's College, London: he is developing an emulator for the Ferranti Mercury.

We are keen to recruit more people to this activity. Some may be deterred through not knowing how to set about it, so we are considering holding an evening meeting in the late spring or early summer to discuss the best ways of proceeding. This would allow Chris Burton and Peter Onion will talk about the problems they have encountered (and surmounted) in developing their emulators, and would provide would-be emulator designers with an opportunity for informal discussions about their ambitions and difficulties.

Pegasus

John Cooper, Chairman

We entered the Pegasus restoration project for the BCS Technical Award last year. During July the BCS Assessment Panel came to see the results of our work, and interviewed some of the working party. We also provided them with our documentation of the progress of the project.

During the Evening Reception following the Open Day, it was announced that the BCS had given the Pegasus project a Commendation. It was the first time the Society had done this—they felt it was more appropriate than a Technical Award. Subsequently, Ewart Willey handed a plaque to commemorate the achievement to the Director of the Science Museum, Dr Cossons.

I'd like to take this opportunity to thank everyone who has taken part in the project, especially those who attended the Assessment Panel meeting.

The machine continues to operate satisfactorily. We are now able to achieve operation on 10% margins on three voltages (−150, 200 and 300).

We have started to repair the broken packages that accumulated during the restoration period. It has been quite stimulating to find how recollections of the old technology came flooding back during this work.

There is now a good chance that the Pegasus will be put on public display in the near future.

S-100 bus

Robin Shirley, Chairman

The main event since the last report has been the donation by Longfield School, Kent, of a complete working Altair 8800b microcomputer, plus parts and spares from several others. I have described the historical significance of this machine in my article 'Altair and After' (see page 23).

As well as system units and spares, the donation includes 8 inch floppy drive units, hard disc controllers and a hard disc unit (all of which are in separate external cabinets). Most of this equipment belongs to the less interesting (but more usable) second generation 8800b series, built after the original Altair manufacturer MITS was bought up by Pertec.

The one exception is a floppy drive unit of original 1975 MITS manufacture, built more like an engineering prototype than something actually sold to users, as those who saw it at the Open Day can attest. The machines were used and upgraded at the school over a number of years, and have relatively recent 64Kb Comart DRAM boards and Soroc VDUs.

The Pertec Altair hard disc unit is of the removable cartridge type used in minicomputers of that time—very solid and heavy. It is not yet running. The floppy disc software uses 16-sector hard sectorised 8 inch discs, and although it has booted successfully under its proprietary Altair Basic operating system (not standard CP/M), we are cautious about going further until we have found formatting and disc copying utilities for this unusual format, so that we can back up the master discs. A copy of Altair CP/M on regular soft sectorised 8 inch discs would be useful—can anyone help?

I have also had a visit from Emmanuel Roche, who has been active in starting up a CP/M Plus ProGroup—a successor to the CP/M User

Group — and is producing a monthly journal. This has already reached its fifth issue (numbered 4, the first issue having been numbered 0!). This contains much material of general historical interest - for example issue 1 is devoted to Alan Turing, including a reprint of the original 1936 Turing machine paper, and issue 2 contains a reprint of von Neumann's 1945 EDVAC report.

Readers who would like to obtain copies of this publication should write to Emmanuel Roche, 8 Rue Herluison, 10000 Troyes, France. I do not know how much each issue costs, but as Emmanuel is a student and is bearing the production cost himself, it would be appropriate to offer him something.

Emmanuel has also preserved the complete CP/M User Group Software Library (on 8 inch SSSD discs), a total of some 120Mb. I hope to have this on PC-compatible media before long.

Forthcoming events

3 February 1993 In steam day

25 February 1993 Evening meeting

3 March 1993 In steam day

25 March 1993 Evening meeting

7 April 1993 In steam day

29 April 1993 Evening meeting

5 May 1993 In steam day

20 May 1993 Seminar on NPL and ACE

24 June 1993 Seminar on restoration of historic computers

In Steam Days start at 10 am and finish at 5 pm. Members are requested to let the secretary know before coming, particularly if bringing visitors. Contact him on 071-938 8196.

Members will be notified about the contents of the remaining evening meetings once the Committee has finalised the 1993 programme. All the evening meetings take place in the Science Museum Lecture Theatre and start at 5.30pm.

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 an annual subscription covering four issues.

Editor – Nicholas Enticknap

Typesetting – Adrian Johnstone and Nicholas Enticknap

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** Manager, Historic Machines Programme, The Science Museum, Exhibition Road, London SW7 2DD. Tel: 071-938 8196

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, Pegasus Working Party **John Cooper MBCS**, 4 Tower Road, Belvedere Road, Kent DA17 6HX. Tel: 03224 30742

Chairman, Elliott 803 Working Party **John Sinclair**, 9 Plummers Lane, Haynes, Bedford MK45 3PL. Tel: 02306 6403

Chairman, Elliott 401 Working Party **Chris Burton**, 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

Archivist **Harold Gearing**, 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