

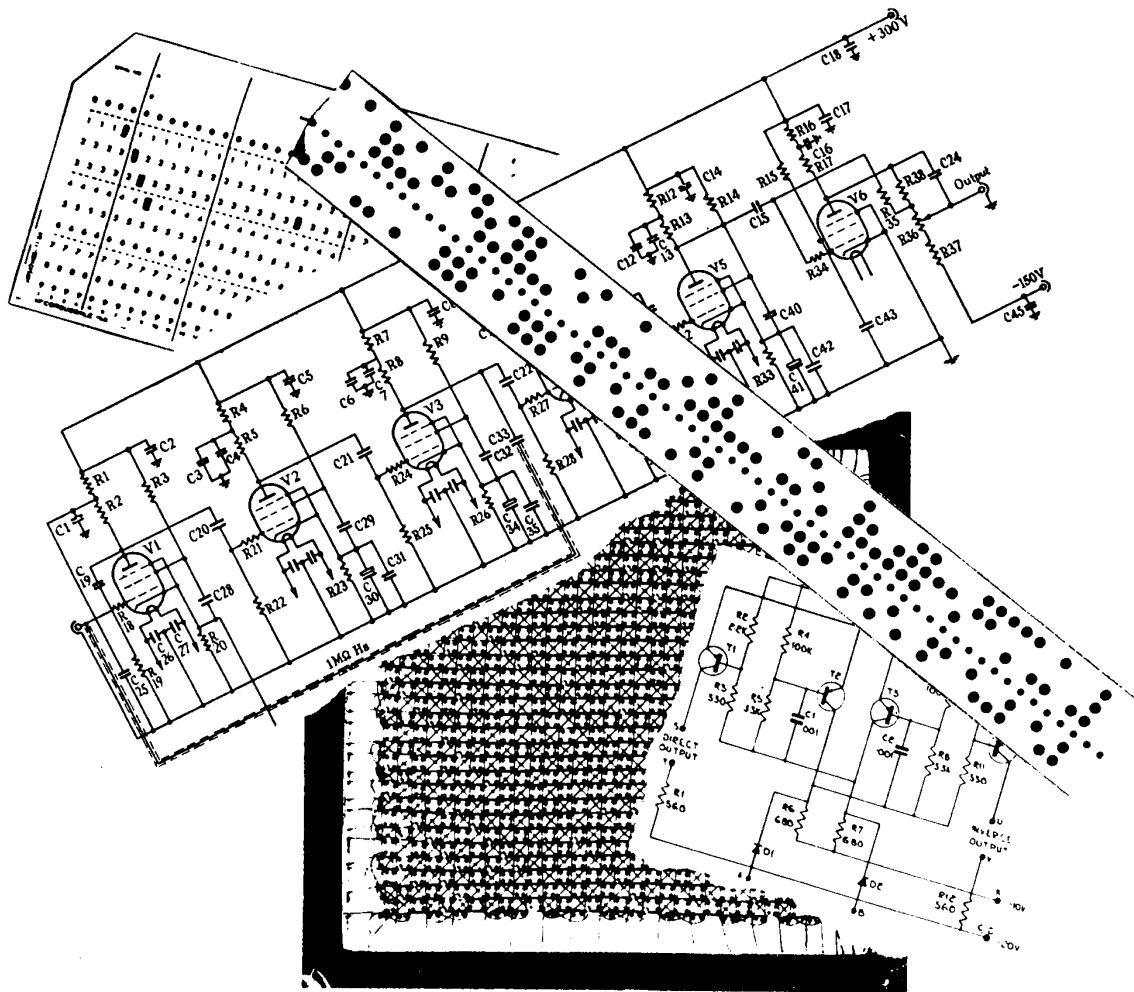
Volume 1 Number 2

Autumn 1990

Computer

RESURRECTION

The Bulletin of the Computer Conservation Society



The British
Computer
Society

&

The National Museum of Science & Industry
Science Museum

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

- o Promote the conservation of historic computers
- o Develop awareness of the importance of historic computers
- o 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

Volume 1 Number 2

Autumn 1990

Contents

Comment	
<i>Nicholas Enticknap, Editor</i>	2
Society news	
<i>Tony Sale, Secretary</i>	3
The thinking behind EDSAC	
<i>Maurice Wilkes</i>	4
Designing a computer for data processing	
<i>John Pinkerton</i>	10
From Cathode Ray Tube to Ferranti Mark I	
<i>Tom Kilburn</i>	16
Evolution of the ACE	
<i>Donald Davies</i>	21
ACE hardware	
<i>David Clayden</i>	25
Working party reports	27
Letters to the editor	31
Skills and Objects Wanted	32
Forthcoming Events	32

Comment

Nicholas Enticknap, Editor

Welcome to the second issue of Resurrection. We were very pleased to receive so many enthusiastic comments about the first, and hope you enjoy this one just as much.

The majority of this issue is devoted to the 24th May seminar, which took place immediately after the Society's first AGM. The unique event allowed members to hear five of the most prominent and influential British computer pioneers of the 1940s tell in their own words the thinking behind the design decisions that resulted in EDSAC, LEO I, the Manchester University Mark I and the Pilot ACE.

Those words are recorded for posterity in the following pages, thanks to the efforts of Pat Woodroffe and Roger Johnson. Pat laboured mightily to reproduce the words for four of the five talks, while Roger organised the transcript of Maurice Wilkes' presentation.

The Society has also recorded the event itself for posterity on film. Dan Hayton and Tony Sale have devoted much time and creative endeavour to producing a video of the event, and this is now available from the Society (details can be found in the Society News section)

The other major Society event since the last issue was the June meeting, which saw the formal birth of the Software Working Party. This is reported in the Working Parties section, which we plan to make a regular feature from now on.

Much of the first year's work of the Society has been taken up with resolving issues such as objectives and methods of procedure, a vital preliminary before restoration work can start. A lot of other unspectacular but vital work has also been done by the Working Parties, such as making hardware inventories and cataloguing software. The greatest progress in actual restoration work has been achieved by the Pegasus and DEC working parties. Our Pegasus is now close to being operational. With the Elliott 803 and Totalisator Working Parties progress has been slower. There is an urgent need for more volunteers to help in their work, especially for people who live or work within easy reach of the Science Museum.

Society news

Tony Sale, Secretary

AGM

The first Annual General Meeting of the Society took place at the Science Museum on 24 May 1990. In light of the Society's extreme youth and with the seminar on "Design Decisions on Early Computers" due to follow immediately afterwards, the proceedings took little time. The meeting formally adopted the Society's constitution, confirmed the existing officers and committee members in their posts, and appointed Frank Hewlett as the Society's auditor.

In his chairman's address, Ewart Willey had plenty of progress to report. In just a few months, the Society had attracted more than 170 members and had set up seven working parties. Treasurer Roger Johnson told the meeting that the Society's finances were in good shape, thanks largely to membership fees and donations from five corporate subscribers.

Videos

A videotape of the historic seminar on 24th May is now available from the Society. Created by Dan Hayton and Tony Sale, "CCS History of Computing Video Archive Number 1" is a full three hour recording of the event. It costs £14.95, or \$29.95 in the USA.

A much shorter tape showing highlights of the event is also being produced for promotional purposes. Interested readers can purchase this one, too, for £9.95. Enquiries to Tony Sale at the Science Museum.

Open Day

We are pleased to report that Computer Weekly has agreed to sponsor our Open Day on 11th October. The newspaper has been a supporter of the Society ever since our formation, and gave us two-page spread coverage of the launch meeting in November 1989.

Totalisator Working Party

We regret to report that John Holden has had to relinquish the chairmanship of this Working Party for personal reasons. We now need a new chairman - will anyone willing to undertake this work please contact Tony Sale at the Science Museum.

Corporate members

We are pleased to report that Allied Business Systems has signed up as our fifth corporate member, joining Bull, DEC, ICL and Unisys.

The thinking behind EDSAC

Maurice Wilkes

When I came back from the United States in August or September 1946, I found myself in charge of the Computer Lab, then called the Mathematical Lab, in Cambridge. Written into my terms of reference were to develop methods of computing and computing machinery, and the purse strings at that time were relatively open. The University like other institutions wanted to get things going again.

So I didn't have to ask anybody "could I build a computer please". I didn't have to put in any proposal. I didn't have to arrange any budget. I was in charge and I could go ahead. The times were extremely abnormal.

There were no staff and so I had to build it up. I was the only one who knew anything about building computers, so if I said "you build a computer this way" they said "yes, right, that's the way". If the first person pronoun comes into this paper as in "I decided this and I decided that" that in fact was the case because of this very anomalous situation.

When we started to do the software and so, on the situation was entirely different. The students had grown up in the lab then, David Wheeler, in particular, and there was a much more normal situation. Bill Renwick joined me in March 1947. It was some time before he got up to speed because I had to tell him what computers were. He was involved of course in some of the details of the implementation.

I shall discuss design decisions under three headings. There is the policy, that is decisions about how to take decisions. There are architectural decisions and there are implementation decisions.

I was perfectly clear in my own mind about policy. It was not a project to build a computer only. It was a project to build a computer, to learn how to use it and then to solve some problems.

The machine was to be simple with no frills, except that it was to be comfortable to use. I didn't want it to be the sort of machine that the user had to know about pulses inside it or timings and so on. There was to be no attempt fully to exploit the technology. Providing it would run and do programs that was enough.

Moreover, there was no attempt to optimise the implementation. There was a rule that if something was made to work, a circuit or some part of the machine, we would accept it provided there was adequate margin: a good margin was insisted upon. If you accepted it there was no attempt to make something that would be more economical or more elegant.

There was no real decision to be taken about the memory. The only memory that you felt you could go ahead with was the mercury memory. Not that anybody ever made one but it all seemed good classical physics - no electrons or anything! I was aware that a number of projects for research in memory, particularly the one that F C Williams and Tom Kilburn were concerned with at TRE and later at Manchester University. But my project was not to be a project on memory development.

Mercury memory was a serial memory and that seemed to lead to a serial machine. Anyway I was under the impression that a serial machine was a lot simpler than a parallel machine. I think I was wrong in that, but that was what I thought.

It was to be a fixed point machine. There were floating point relay machines but none of us could face it electronically. Even the people who were trying to build a very much higher class machine than mine didn't build floating point either until much later in the game.

One question was how negative numbers were to be represented. There were three possibilities. The one we went for was the one that is popular now, namely twos-complement. Some people favoured sign plus modulus and some people even favoured the ones complement system in which you got the negative by replacing each digit in it by its inverse - 0 for 1, 1 for 0. Both those two latter systems are a bit awkward, for they involve two zeros, plus-zero and minus-zero.

How long was the word to be? Well I reasoned as follows: 10 digit decimal accuracy was about what you want. That is 34 binary digits. Add one for a sign, which makes 35, and another one for a space between words giving you 36. This was really rather nice because half of 36 is 18: knock one off for a space, leaving you 17, and 17 bits it turned out were just about right for an instruction, allowing two instructions per 36-bit word.

I knew we were going to be short of memory, and I thought as we have got to have circuits for dealing with half words for instructions, we might as well make half words available for numbers. I thought that many numbers could be represented to half word accuracy and this would make the best use of the small amount of memory we were going to have.

Memory was addressed by half words so if you wanted a full word, a long word, its address was even - none of this left and right business that appeared in some machines.

Another decision was the position of the binary point. If all numbers are fractions in a twos-complement machine, numbers lie in the range $-1 < x < 1$. That is to say, -1 is in the table but +1 is not. I pondered over this and thought

it might be rather awkward not to be able to have the number one, a number which comes frequently into calculations!

So there was a scheme for a time of having numbers in the range of $-2 < x < 2$. I am not quite sure how it happened but we didn't use it. I used to like to say that when the arithmetic unit was built it turned out that the numbers were in the right range $-1 < x < 1$, and there may be some truth in that. Anyway fairly late in the design process we realised that it was better to have $-1 < x < 1$.

I decided to go for paper tape for input and a directly coupled teleprinter for the output. The alternative was punched card. I think I felt that punched card equipment would be more complex - there was more to build. Also it would have involved negotiations with manufacturers of such machinery whereas anybody could buy teleprinter equipment. There were one or two factors favouring the teleprinter equipment. For example, on input one of the bits was reversed so that blank tape didn't mean anything.

The single address instruction set followed naturally from the serial architecture. It is easy in a serial machine to make the accumulator double length. That is to say if two numbers were multiplied together then all the bits, the double length result, went into the accumulator and was added to what was there.

The way the programmer did it was to load the register called the multiply register with the multiplier, and then use the multiply instruction to bring from memory another number, and then the multiplication was performed by hard-wired logic. The beauty of this arrangement is that you can accumulate a scalar product with complete accuracy and do the rounding off at the very end.

Having had some exposure to numerical analysis I realised that the most important thing you ever do in numerical analysis is rounding off, and therefore we did it properly. We had an instruction for rounding off. It added one in the base in the accumulator just to the right of the least significant bit.

We did cut a corner over the shift instruction. The obvious way to do the shift instruction was to have an op-code that said left or right, and then in the address part of the instruction have a number. For example, if the number 5 came then there would be a shift of five places.

Now the hardware was designed so that the shifts were of one place at a time, and to get a shift of five would immediately be five shifts and a counter, to count that the right number, in this case five, had been done. I didn't want to build a counter. I wanted to save equipment. I therefore decided that instead of having the number five the amount of shifting would be determined by the position of a one in the address field. If that one was in the

least significant place we got a shift of one place. If that one were in the next most significant place there was a shift of two, if three along a shift of three and so on.

That simplified the hardware. It meant you couldn't do a shift of more than 12 with a single instruction. If you wanted to do more than 12 you had to have two instructions.

I believe this may have saved us a few weeks in the construction. It was horribly inelegant and an absolute pest but there it was. Looking back, I like to think perhaps I ought to have spent a little more time seeing whether there wasn't some cute way of doing it, but in fact I didn't and that was how it was done and that was why it was done.

There was another thing some people thought was very peculiar. Many machines that had been designed at the time on these principles had an add instruction which would add a number into the accumulator, adding it to what was there. They also had a clear and add instruction, that would first clear the accumulator and then insert the new number, and they would have a subtract and a clear and subtract.

I thought you could save an instruction by doing it another way. This was to have two transfer operations that would copy the number in the accumulator to memory. One of those would leave the number where it was and the other would clear the accumulator. I got a good deal of flak for this, from Christopher Strachey in particular.

Another thing about the instruction set was that I do think we did the input and output rather elegantly. The input operation would read a single row of holes on the tape and put the resulting five bits into memory. The output instruction would take five bits from memory and send them to the teleprinter where they would be printed as a character. We might have gone through the accumulator instead of through the memory but it wouldn't have made all that difference. I think it was a nice, neat way of doing it.

Now I come to implementation decisions. It had to be a mercury memory. I had an enormous stroke of luck in that I met Tommy Gold, who gave me a design for a mercury tank which he had evolved or known about when he was working with the Admiralty during the war. It had not been used for digital memory but for something else in radar, but at any rate there it was and it worked. I built a single tank to his specifications. (We used to say tank instead of tube for these mercury tanks, in order to avoid confusion with vacuum tubes.)

We had 32 tanks in the end and they all had to be exactly the same length to within a very close tolerance. Moreover the crystals had to be aligned very accurately. The crystal was about half an inch across and that is an enormous

number of wavelengths of sound in mercury at 13.5 megacycles. The beam that comes off the crystal is a very, very sharp one and it has to hit the crystal at the other end. The walls of the tube are not there to guide the sound as it goes down, but to keep the mercury in. The beam of sound is very sharply pointed and doesn't touch the walls at all. The accuracy was of the order of a few minutes of arc.

One way of doing this was to make the tanks with adjustment screws, so that you could adjust the tanks to the same length and adjust the orientation of the crystal so that the beam hit it off. I didn't want to do it that way.

I thought it would be very much nicer to build these tanks in batteries of 16, to design the battery carefully, have it made with sufficient precision so that when it was bolted up solid everything was in line and the crystals were the correct distance apart.

I was very fortunate in that the superintendent of the central workshops in Cambridge in the Engineering Department liked that sort of job. It was a challenge but not a desperate challenge for him to meet the tolerances. Nevertheless I must say I was very relieved when we had a battery and filled it and all the tanks worked. I suppose there was a certain amount of risk but I am sure it was a good way to do it.

The registers in the arithmetic unit - the accumulator and the multiply register and the registers for holding instructions while they were being expected - were all short mercury tanks. The main tanks were about five feet long, the short ones were an inch or so long.

Then there was the question to be decided about logic: how the switching was to be done. An obvious way was to use pentodes, and to use the control grid and suppression grid. I did a few experiments with this and it didn't seem very nice.

The suppression grid needed a much bigger swing than the control grid. That made it horribly unbalanced. It was all, I thought, pretty messy. So we used diode switches in the thermionic valves.

They were driven by cathode followers, there was a pull-up resistor, and the output was put into an amplifier. I was very proud of my discovery, after quite a lot of experimental work, that the ideal sort of amplifier for that purpose was a cathode-coupled amplifier.

The machine was AC-coupled and extensive use was made of DC restoring diodes to keep the level baseline in a fixed position. I soon discovered that there could be no compromise about that: we had to put these things everywhere and make sure they worked.

It was also necessary to make sure that in between the pulses the wave form came down to zero. It was no good any one pulse running into another in our design. The pulse had to come down and remain down for an appreciable time at zero before it went up for the next pulse

Now this was much helped by a decision I ought to have mentioned right at the beginning. In order to make the implementation job easier and therefore reach earlier the point at which we could start programming, I decided that it would run at half a megacycle per second rather than one megacycle per second. Any electronic engineer worth his salt would take the challenge of working at at least one megacycle and all my colleagues were working on this kind of machine, but I think we bought a great deal by making that decision. It certainly made this question of getting nice clean boxes with proper spaces in between very much easier.

On the mechanical design, let me tell you why we didn't use 19" racks. I never liked 19" racks. Why something that appeals to people who design telephone exchanges should have exercised such a grip on the electronics industry, I have always failed to understand! So we had a long chassis. They were bent in a strange way, but they were very good, and made everything very accessible.

Each chassis had a heater transformer on it. The chassis were wired in, not plugged in. I suppose I was a bit scared of plugs which in those days had a bad name, perhaps a little unfairly. Also I was afraid of the capacity that there would be on the signal wires and in fact the chassis were soldered in. The signal wires that ran around the machine were just open wires that looped across and were soldered on to the tags on the chassis. All very simple and quite effective.

Designing a computer for data processing

John Pinkerton

Some of the points about LEO are very elementary and obvious. It was enormous and of course it was very slow. It worked at the same speed as EDSAC because it used essentially the same logic and the same circuitry; but it was also very unreliable.

We were very much indebted to Maurice Wilkes but we had to recognise that when you did clerical work there were some differences. One of the major differences was that there was an enormous amount of input and output to do and relatively little calculation. This is of course something which is well recognised today but it wasn't quite so fully appreciated at the beginning; therefore the problems of input and output loomed very large and so did unreliability. I'll deal with them separately because the two aspects were dealt with in different ways.

The reason why it was slow on clerical work was, first of all, there was this very large amount of input and output and secondly, because of the conversion into and out of binary that you had to do. If that were done by program instructions in the manner that was being developed at Cambridge it would have taken about 90% of the total time; and that was much longer than we were prepared to spend on it. Also there was a tremendous lot of input to feed and a lot of output to record and we didn't want to slow the machine down by having to wait for mechanisms to do things. So that was our initial conception of the problem.

We expected that there would be considerable unreliability from the use of valves. We found that valves were unreliable from two causes. One is that they suffered from inter-electrode insulation troubles and so on, but I think the major cause of trouble is that their emission gradually declined and after a while they lost enough emission so that they didn't function in the circuits properly, and you had to throw them away.

Now let's focus on the first approach to the problem of input and output. We would have multiple input and multiple output channels and they would be buffered so that instead of feeding one character at a time we would feed a whole group of numbers at one time; we would hold them in a long delay line which would act as an input buffer; in fact we had double buffers on the input in order to give the possibility of filling one buffer from whatever kind of reading device we were going to use while the other one was ready, full, to be pulled into the machine in a hurry.

We also looked at this conversion and re-conversion problem and we were frightened by it frankly. We didn't see how to do it and we thought the answer was to do it by some external means in the input and output channel.

Now it happened that at this time Lyons had just bought a telephone exchange from STC. It wasn't connected to GPO lines at the time but it was a very good internal automatic telephone exchange. So STC had a very good reputation with Lyons. So Lyons got in touch with STC and said "Do you think you could find some way in your research of constructing external converters and re-converters to change numbers from decimal (or Sterling incidentally) into binary and vice-versa?". Within a matter of weeks they had worked out how to design the logic of converters and re-converters using a kind of circuit technology known as the gas trigger tube.

The buffers in each of the input channels were all valve circuits and mercury delay lines. They had little short tubes which assembled individual numbers and put them into position, one after the other until they filled it up. The same thing on the output side; some of the buffers were doubled on the input though not on the output. The logic allowed four input channels and four output channels, though we only ever built three in and two out. Between the things which were going to read the data to be recorded on the magnetic tape were gas tube converters and on the output side were more gas tube re-converters.

The design seemed satisfactory but it had a rather sad history. This was before transistors were invented and STC were interested in gas trigger tubes for use in telegraph repeaters. They had built some telegraph repeaters which worked quite well.

The trouble was they suffered from a certain form of unreliability which meant that if you made a trigger - a flip-flop - with two trigger tubes coupled together then occasionally they both came on together. When that happened there was no way of getting it into a correct state; you had to pull one of the tubes out. You could see that this had happened because both tubes were alight: that gave the game away. We also used multi-cathode tubes which were similar to the Ericsson decatron but used more current so they would drive a gate directly. They too had some trouble with the physics because occasionally you got two cathodes up together and two glows stepping round at once which of course didn't have any sensible meaning.

There was a tremendous number of components associated with the transports. STC dreamt up the idea of having what they called swinging gates, which were boards sticking out backwards to hold the very large number of components. The tape deck is of very considerable interest; it used 1/4" tape pushed into a box without being wound up. It was an endless loop of tape; and I think Maurice Wilkes once said "the tape couldn't get knotted up because you can't tie a knot in two dimensions" - because the tape can't get past itself. It can all be in the outer loop or a lot of it can be in the inner loop.

There were two points at which the tape could be driven and two heads, one for reading and one for writing. The idea of that was that you could have somebody recording material from a teleprinter a character at a time and somebody else checking what they had just recorded a little bit further down the same tape, which was then fed past the read head. So as long as they didn't overtake the recorder they could check at the same time as recording. So Standard Telephones, besides doing the converters and re-converters, invented circuitry for data checking.

As I have said this did not work out satisfactorily simply because there were quite a lot of things wrong. The electro-mechanical design wasn't good but I think it could probably have been corrected; but I don't think they could have got the physics of the gas-tubes right. Anyway they were overtaken by transistors; transistors made them rather unattractive.

After some two years of trying to get the tape system to work Lyons became rather impatient and told STC to take it away, but not until we could think out some way of substituting for the gas-tubes to do the conversion and re-conversion. Now very fortunately somebody had put on the market a germanium diode instead of the hot cathode diodes which we had used, like EDSAC, for the logic. They were very small, about the size of a 1/4W resistor; so you could contemplate using a very large number of them.

We found you could do a conversion and re-conversion instruction using a diode matrix to store the equivalent numbers. So if you stored the binary equivalent of 10 and were in a position to shift it one, two, or three places you could also generate the binary versions of 20, 40 and 80. We simply used the same control circuitry that was there, which Maurice Wilkes had provided very kindly for doing multiplication, to do conversion by simply looking at the successive values of binary coded decimal numbers and adding the difference into the accumulator. This is a process which was dreamt up by Lennox and it worked fine. I think we used something like a thousand diodes to make the diode matrices. Re-conversion is obviously an analogue of division and provided you store the equivalents of the Sterling numbers you can do Sterling conversion in the same way. So we had alternative instructions to do Sterling or decimal conversion.

That was a completely satisfactory answer but we still had to deal with the question of the great volumes of data that needed to be fed in, and the great volumes of results. So we decided to be very conservative rather than innovative as we thought in the case of the STC tape-decks. (I'm afraid we took everything that STC said they would do on trust which was obviously rather silly; we should have been more sceptical). So we decided we would use punched card equipment. We went to the British Tabulating Machine Company, and we bought a couple of tabulators of a rather peculiar design which had what they called "pence bars" in every position. That is to say they

would print numbers from 0 to 11 in every position across the line. We coupled those to the machine and we could print using a 35-bit word on 35 columns at a time. We had a switchover so you could then print on another 35 columns, so you could either duplicate what was on one side of the paper on the other, or you could print something differently before you move the paper up. So in fact we had a very powerful system of input and output using punched cards and tabulators, but we retained paper tape.

The Lyons idea was: there were three kinds of input, what they called "current" data (what had happened this week if you like on a payroll), "brought-forward" data (which was what had happened last week which had to be fed in again) and "permanent" data which was data that didn't change very much and it was convenient to have it on a separate pack of cards. So you had the current data on paper tape, you had the brought forward-data on punched cards and the permanent data on punched cards. All these three channels of data had to be kept in step so that the data for the same individual on the payroll was coordinated.

So we went back to classical peripherals and we abandoned the gas tubes. After we had done that the buffering arrangements stayed the same but the converters had disappeared because they were now dealt with inside the machine by convert and reconvert instructions. We had two paper tape readers and three punched card readers any two of which we could use at a time; they were switched through two channels. So normally one of them would be reading the brought-forward data and one would be reading the permanent data like people's names and addresses if they were involved or other details about people that didn't change.

On the output side there were two line printers one of which could be being maintained, and two card punches, one of which again could be being maintained. That was because we had this phobia about the reliability of input and output machinery. It turned out actually to be quite a good idea because we found that as soon as the punched card machine had been maintained it was more unreliable than it had been before, but in the long term it was better to have it maintained than not.

We got to the stage where you could do payroll more or less at the cadence speed of the printer which was 100 lines per minute. We devoted two lines to each man so we were doing 50 men's payrolls per minute; which considering the machine was so slow was not bad. The rhythm of that machine running is something I can still hear in my head - clanking away day and night turning out payroll and other jobs.

Now the reliability question. There are certain things you can do about reliability in the design, there are certain things you can do about it in

maintenance to offset the inherent unreliability and there are certain things you can do about it in operating.

I think all those three things have to be considered. In the design you underran all your components, and where you needed a 1/4 watt resistor you used a 1/2 watt, and that sort of thing. We underran the valve heaters; the cathodes would last longer if that happened. We brought the heaters up very slowly and took them down very slowly over about a minute with a special kind of variable transformer and generally speaking everything was kept under as close a control as possible. We used anti-parasitic grid stoppers in every valve and we were really a bit fanatical about interference. We did actually suspect that we were going to get a lot of trouble from interference from the punched card equipment which was full of contacts which opened and closed large currents in inductive circuits. So every contact had to have spark suppression put into it to make sure there wasn't any electrical interference.

We had a certain amount of trouble with electrical interference from the fluorescent lights at the start, because in an ultrasonic store you have signals of about one millivolt and if you're not careful you can get interference into the store. But we were a bit obsessive about it. Our attitude was: it's cheaper to take precautions than to find out whether they were actually necessary or not.

We took a lot of steps to check that valves were in good shape. We maintained them by numbering every valve. At regular intervals each valve was taken out and had its mutual conductance and inter-electrode insulation checked, and various other things were done too. In operating again you had to work on the assumption that the machine was probably going to break down. It had a mean time between failures of, I suppose, between six and 10 hours that would be my mental recollection of it. That meant that you had a fairly good sporting chance of having a failure while doing your job. So you couldn't afford to run a job that might take say two hours without having some restart facilities in the middle. The phrase "checkpointing" hadn't been invented, but that's what we did. We took departmental totals or other totals out at regular intervals, so that if the machine showed that there had been an error or it had stopped or broken down in any way, then instead of having to go back to the beginning you went back to the last point at which numbers had been discharged. That was left to the operator as there was no operating system, of course, possible. In fact nobody had even heard of an operating system; such a thing wasn't spoken about.

LEO was very unreliable but it was effective because we took these steps to offset the unreliability and to work through it so to speak. I think it laid the foundations for data processing, I don't think that's an unreasonable claim.

The fact that other people did it the same way without knowing the way we'd done it doesn't detract from the value of what was done.

A lot of people came to see what was done and it's impossible to say now how much inspiration it gave. I think probably the biggest inspiration was that people saw that Lyons had managed to do it at all. They probably didn't pay too much attention to the way it was done but the fact that Lyons were doing it was a great stimulus to other people to go away and do something similar.

From Cathode Ray Tube to Ferranti Mark I

Tom Kilburn

Before I start I would like to acknowledge the contributions of Freddy Williams, Geoff Tootill, who was a colleague of ours from TRE, Alec Robinson who worked on the multiplier, Dave Wood who worked with me on things like the main machine and the B-tube and so on, and Tommy Thomas who worked on the magnetic drum.

Freddy Williams went to the States about 1945 and 1946 to contribute to a set of radar books that were being written. He saw at Bell Labs some experiments on the cancellation of ground echoes in radar which involved moving signals from a cathode ray tube (obtained by having a metal plate on the face of the cathode ray tube) from one tube to another tube and then back again, as a method of storage. Freddy came back to TRE about August/September time and in '46 he started to set up this sort of system with a view to trying to store digital patterns. By December he had stored one digit at TRE!

Freddy Williams' idea was that one type of pulse which he called an anticipation pulse told a circuit or an individual that there was going to be this gap in the trace. So this was called the "anticipation pulse system" and, as I say, it was developed in December 1946.

By that time I had been in his group for four years and he invited me to accompany him to Manchester. I stayed as a member of TRE but on outside duty at Manchester in order that I could draw stores from TRE!

Now for three months or so, largely at the beginning of 1947, I did experiments varying the speed of the trace and the focus of the trace and so on. I came to the conclusion that another pulse, which didn't play any part in the original anticipation pulse, was much more useful than the anticipation pulse itself; and so sometime in March 1947 I convinced him that we could drop the anticipation pulse and use the positive pulse. The reason the positive pulse is better is that it's more uniform over the whole face of the cathode ray tube; and also the distinction in the anticipation pulse between a one and a nought is between a negative pulse and zero whereas the distinction in a positive pulse is between a positive pulse and a negative pulse, and any engineer knows that that gives you a better definition.

So we switched to that system and this is offered in a lot of schemes including both the dot/dash system of storage and the defocus/focus system of storage, which is the one that actually went into the Ferranti machine. The original pulse that we started off with was abandoned.

The whole thing was enclosed in a metal box, because we found in early experiments that if you left it in the wooden crate that it came in and just put a metal plate over it and fastened that to the amplifier, anybody riding down the road outside the lab on a motor-bike would fill the whole thing with random noughts and ones so it was absolutely essential to have the screen. The signal was of a few microamps and the amplifier provided a gain about a thousand giving an output signal of, say, 10 or 20 volts. The finished result was actually a rather elegant piece of engineering by Ferranti.

Of course if you improved the cathode ray tube you could get improved performance, and in particular I remember a few things. If you used soda glass for the cathode ray tube then nasty charges built up on the surface of the glass which ruined the system. So lead glass was an advantage. Also it was necessary to improve the focus in order to get any more digits on the tube, and so we asked GEC to work on this.

It was also necessary that the screen on which the charges on the inside of the glass were deposited was pure, and IBM, who licensed this type of store to start off their 700 series, told me that although they'd moved their cathode ray tube factory up the Hudson to Poughkeepsie to get rid of the dirt, they were actually troubled by pollen. Of course, as you made the focus much, much better, the necessity for a purer and purer screen became more important. So GEC did quite a lot of work on this cathode ray tube business. Eventually we stored a maximum of 2560 digits on a Cathode Ray Tube. This was using the defocus/focus system where large dots are a defocussed spot representing a one (a negative pulse) and small spots are focussed spots representing a nought.

When we went to the Ferranti machine, for increased reliability we halved the number of digits again, and went back to 32 by 40 in order to make sure that the thing worked really well. If I'd been able to sit with it night and day when I could have kept the larger screen working but of course you can't sit with it night and day.

The cathode ray tube was housed in 19 inch racks, because we'd worked with them for four years at TRE in radar experiments. Together with a keyboard, it formed a primitive form of VDU which was a method of input to our computer.

The first small machine worked in June 1948. It's interesting why we built that small machine. We had a cathode ray tube which would store patterns on the CRT store over long periods but it wasn't really proof that the cathode ray tube system would work in a computer, because if at very high speed you write noughts over ones, or ones over noughts, which is what you are doing in a computer constantly, the signals you get from the screen are not balanced and the base line starts heaving up and down. We'd never seen this

heaving up and down because we'd never fed it quickly enough but we surmised that it would be there and indeed it was.

So I decided to design some gear which would test this, but after a few weeks (actually I was travelling into Yorkshire at the time in that awful winter of 1947 and I did a lot of design on the train) one of the conclusions I came to was that the only way to test whether the cathode ray tube system would work in a computer was, in fact, to build a computer. So I designed the smallest computer which was a true computer (that is a stored program computer) which I could devise, and we ended up with a one-tube, 32-lines, 8-digit machine. The signals did in fact heave up and down and the design of the amplifier and the clamping system to deal with that was quite an interesting exercise.

Now when that machine worked I had, so to speak, completed my duties because I'd been released to work on this system by TRE. And so I wrote up my thesis for a Ph.D. on this system, and decided to take a three month's holiday before returning to TRE in December - this was June of course. But this didn't work out, and the reason was that we had a number of visitors.

The most important of these were Sir Henry Tizard, who was the chairman of the Advisory Committee on Scientific Policy, and Sir Ben Lockspeiser, who was the Chief Scientist at the Ministry of Supply. They were very interested in the machine. The upshot was that Ben Lockspeiser arranged a contract which was given to Ferrantis to provide a copy of a machine which the University would design. So I spent the best part of the next 18 months designing that machine.

However, the most interesting visitor was Wilkie Wilkinson from NPL, who had in fact been a tutor of mine at Cambridge. He came up one day and he was thrilled to see this machine working because he'd been landed with the job of trying to make the NPL machine work. I remember during our conversation we discussed the length that a word should have, and we came up with the fact that 32 digits were a bit skimpy. I wanted 40 digits because that would give me two 20 digit instructions per line. 20 digits seemed quite generous at the time, but it was a fortunate decision because the invention of the B-tube meant that we needed three of the instruction digits for the index register of the B-tube.

That was one thing we discussed; the other thing we discussed was single and double length instructions, and we came to an agreement that a double length instruction, that is one with two addresses, was worth 1.3 times a single address instruction. This confirmed me in the view that we should indeed stick to a single address instruction. Going to a double length instruction was not worth it on any grounds and would have ruled out having two instructions per word in any case.

From about November when the contract with Ferranti was signed we went on to design substantially larger machines in which there were two main features which stand out above all the others. One was the invention of the B-line, which is quite an interesting story.

Most of the people working at Manchester were thinking of different structures. (The people I'm considering are Williams, myself, Newman and Tootill.) Things like two accumulators and so forth. I myself was interested in splitting the instructions from the numbers because this would have allowed me to double the speed of the machine. While extracting a number from one half of the store I could have been getting the next instruction from the other half. But I was advised by all the programmers I spoke to that I couldn't do this. So the next thing I proposed was that we should have a separate accumulator for instructions: would that be all right? The answer was still "no". But you see the kind of thing that was going on; we were trying to think of different structures.

Now out of one conversation between the four people I've mentioned came the index register as being the best new suggestion for the machine; and so we put in the B-tube. It's called the B-tube because we had an accumulator and a control - A and C - and so we called it "B".

The other very interesting feature is that we decided to link the Cathode Ray Tube store to a backing store using a magnetic drum. The magnetic drum was 1024 words and it was a 10 microsecond system. In fact all our work was done at 10 microseconds, and that made it very much easier than the one microsecond that Maurice Wilkes was trying to deal with at the same time. The cathode ray tube store ran at 8.5 microseconds in its early days but when it went into the Ferranti machine, again for practical reasons, we expanded it to 10 microseconds.

Of course the immediate access feature of the cathode ray tube store more than compensated for the slower digit rate. That is to say we could switch to any line within 40 microseconds which couldn't be done with any delay line system. So we came up with much the same speeds of instructions as Maurice Wilkes at Cambridge.

There's an interesting thing about the magnetic drum and the order code. In the order code I introduced a relative control transfer with the idea that if, say, a subroutine were put in a different place in the store from which it had been programmed, the relative control transfer would not need to be changed, because it simply moved you back and forward relative to the instruction which you were at.

The other thing was I put a tag line which we called the 65th line on each of the pages or blocks of words on the magnetic drum. So there were 64 useful lines and a tag line-the 65th line. This was with some vague notion that

programmers would be able to move tracks around on the drum, not necessarily in the fixed order in which they had originally appeared on the drum, and still keep track of them. Now the programmers who used the machine never took advantage of this, but these two sorts of ideas were the seeds of the virtual memory and the one-level store which came along with later machines. So those are of special interest to me.

I think everybody had a lot of discussion about the incidence of the various orders; that is, in particular how many multiplications would there be relative to all the other instructions. We found that 16% of all time was spent on magnetic drum transfer, 28% was spent on multiplication and 56% of time was spent on other orders like addition.

Now what that means is that if we had gone into the Ferranti machine design with a slower multiplier the machine would in fact have been five times slower than was the actual case because what we did was choose the degree of parallel multiplier which made sense. The other orders took 1.2 milliseconds or 9 milliseconds if they were B orders while the multiplier took 2.16 milliseconds. The multiplier was actually a half parallel multiplier - it multiplied 40 digits by 20 and then 40 digits by the remaining 20 - so two stabs were taken to get a 40 x 40 multiplication and that took a total time of 2.16 milliseconds. Quite a surprising thing also is that one in five of all orders obeyed (not written - we're talking about orders obeyed) were multiplications which is really quite startling I always thought.

Evolution of the ACE

Donald Davies

ACE stands for Automatic Computing Engine - Engine by deference to Babbage. Its origins were a report by Turing in 1945 to the Executive Committee of the National Physical Laboratory, which had to approve the project - which they did.

Turing's report, which has recently been reprinted, is a remarkable document, setting out both the reasons for carrying out the project and what Turing believed the machine should be. He describes its use on numerical calculations of interest to scientists and engineers; also on clerical work, though he points out this would require excessive amounts of input and output; and also on logical and combinatorial problems and number theory. Though not mentioned in the report he also considered its use in weather forecasting; I remember discussing it with him. So it had a very wide range of intended uses mainly involving large scale mathematical calculations.

Turing's report doesn't give much detail of the design, but describes in detail the notion of optimum programming. He was already convinced that ACE had to use a delay line store; he hadn't heard of any other possibilities, and with a delay line store you have to wait for the numbers to come round. Optimum programming avoids some of the time wasted in doing this.

The technique involves placing instructions in the delay lines where they will appear at the right time; also judicious use of single, double and quadruple word as well as long (32 word) lines, using great skill in laying out the numbers. If numbers had to be processed sequentially it was arranged that a long delay line would precess, so that every time a number was extracted all the remaining numbers would shift round.

The technique was very successful in avoiding the latency problem but involved an extra stage of mental juggling. After writing the programs you had to decide where to put everything in the delay lines. Clearly this was not a machine designed for people coming in every day and using it in the way we use personal computers now; it was a machine designed for people devoted to programming and getting the maximum speed out of the machine. Of course as soon as we had easily available random access stores it was no longer of interest, so it was a temporary phase in the history of computing, an interesting one and one that did extract useful speed.

It was particularly valuable when one could spend a lot of time optimising the inner loop of a program which will be repeated very many times. This influenced the type of application that the machine executed most efficiently - numerical analysis in general but in particular the manipulation of linear equations and matrices. The result was a vast library of extremely powerful

functions to perform on matrices which turned in the end into an interpretive system by which one could quickly assemble these to do matrix operations with enormous efficiency.

Turing's proposal was in 1945. After that, a group in Mathematics Division kept writing programs to try to improve their efficiency; and meanwhile hardware was being developed by the Post Office Research Station at Dollis Hill. This has an interesting background too.

It seems that during the discussions by the Executive Committee one of the members wanted the computer to be designed and developed because of the requirements of cryptanalysis, and Turing's involvement at Bletchley Park had something to do with it. In fact the people at Dollis Hill who had built Colossus were brought into the hardware team. However this arrangement didn't work well and NPL and later Dollis Hill separated.

Dollis Hill went ahead to build the ACE to NPL's design and it became the MOSAIC, a Ministry of Defence machine. Then Harry Husky joined the team from America where he had been working on ENIAC. He put a lot of energy into the team which had been flagging after the poor beginnings. He decided to build a Test Assembly in order to try out the principles of the machine quickly. That was a very simple design. Turing himself had left by that time to go to Cambridge and later to Manchester.

This was the genesis of the Pilot model. The Test Assembly was never completed; its electronic and mechanical design was rather "hairy" and I sometimes wonder if it would ever have worked. Shortly after Husky left it was decided to close down the work on it.

Still things didn't go well. In 1948 there were two separate teams at NPL; one was the so called Electronics Section and the other in Mathematics Division where they were still writing programs. It must have been very boring writing programs for so long that you could never test. We knew they would be full of bugs but we could never find them.

Finally, after three years of negligible progress, early in 1949 the two teams came together. All the interested people in Maths Division walked out and joined the Electronics section and everybody started building. We forgot about programming and set to with soldering irons. This arrangement worked well.

From that time things moved very fast and the first programs were running haltingly in 1950, and by the end of that year useful programs were running. In 1951 we had a multiplier fitted and a new control system designed by Ted Newman which saved some of the logic. That completed its development, apart from the magnetic drum which came in 1954 and was very important. It was operating as a service from 1952 onwards and was

moved to a new location where it was possibly the first machine selling its services commercially to the aircraft industry among others. I remember doing work for the Safety in Mines Research Board and the Ministry of Transport.

One of the peculiarities of optimum programming is that you have to have a complicated instruction to deal with it. The first field of the instruction is of three digits and gives the next instruction source - that is the tank from which you take the next instruction. Next come the source and destination fields; these are not simply stores but they also carry the operation code with them. So everything about the operation is defined by the source and destination, which is another peculiar property introduced, I think, principally for the Pilot model. Then there are two timing numbers, Wait and Timing, and a Go digit which enables one to stop and step on and which is also used in synchronising input and output.

Transfer lengths are either single word or multiple; these are very valuable as one could transfer the whole of a long tank into another one with one instruction; or one could add a whole series of numbers into a short, double or quadruple tank. So flexible operations which we now take for granted were possible here, albeit in a primitive form.

The Wait number defined the time before the transfer started, and the Timing number determined the end of the transfer which was immediately followed by the next instruction. If the Serial bit was set a single length transfer occurred. So the instruction form was quite complicated: quite powerful if one knew how to use it, but very tricky to do well.

The ingenious method of dealing with initial orders to get the machine started up was similar to that used at Cambridge. The first instruction was always zeros, set by pressing the initial input key. This was interpreted as "take the input from the Hollerith reader and feed it direct to the instruction register". In this way the first line on the card was taken as a command and obeyed immediately. The card, by the way, was punched horizontally, not in the usual way. This gives 32 binary digits across the card.

A word about input and output. We naturally chose Hollerith machines as they were widely used in the Maths division and well understood. It fell to me after Stanley Gill left to carry out the engineering of the machines for input and output, so I had all the problems of noise. These were overcome and they worked well; later on faster machines were installed. Our input speed was 12 32-bit numbers per card at 200 cards per minute and output was at 100 cards per minute. This was much faster than we could have done with paper tape and very soon became highly reliable, as we learnt how to carry out the necessary adjustments.

Cards were in fact used as our first level of store long before we had the drum. For large matrix operations we had to punch out intermediate results, very often rearranging the cards as part of the algorithm and then feeding them back in through the reader. This was the only way of inverting very large matrices.

To summarise, the machine was built to meet the requirement of the Mathematics division, primarily for numerical analysis. It was aimed at the fastest operation that we could get with delay line storage. It went through a number of stages and eventually ACE was built but by that time it was a dinosaur, out of its time. Core store had been invented and I tried unsuccessfully to deflect ACE in that direction but it had too much momentum. There was a three year delay before it started due to organisational problems.

The Pilot model achieved all its objectives. It was fast; it was very simple having only 800 valves. It led to three machines; to Mosaic which I have already mentioned; to DEUCE of which I think 32 were made and very widely used; and to the Bendix G15 made in USA by Husky after he returned. He designed it around a magnetic drum, replacing the tanks. In fact the drum had two heads, reading and writing simultaneously. The G15 was very successful and 400 were made. The 400th was the last one and it was gold-plated. Husky had one in his garage so he must have been the first person to have a personal computer! It gave out a great deal of heat which must have been rather unpleasant but it worked.

ACE hardware

David Clayden

I arrived at the NPL in 1947, having been an engineer at EMI Hayes. My first job was to make the mercury delay line work with some electronics which had been made by a contractor. I would like to describe how mercury delay lines are made.

You start with a straight column of mercury. At the end is a piezo-electric crystal which is driven by an electrode behind it. There is a gap between the electrode and the crystal; this is very important, because if the electrode touches the crystal it damps it enormously. So it is capacitatively coupled to the crystal, while the mercury makes contact with the other face. The whole thing is sealed (you hope) so that the mercury doesn't run round behind the crystal and is held together with bolts. A large voltage at about 15MHz is applied to the electrode, and this is converted to a sonic pulse which travels through the mercury. At the other end the sonic pulse is converted back to a voltage. The total attenuation in the long delay lines is about 60 db (that is about 1000:1).

In a typical ACE circuit there is a long tailed pair - a double triode. The total current is defined because the cathode resistor is returned to -300 volts and the grids sit at about -200 volts so there is 100 volts across the 33k resistor giving us a defined 3 mA, so it doesn't matter what the grid bias is on these valves. As the valves deteriorate they continue to operate until they run into grid current.

Another piece of standard technology used in the ACE was dc coupling. A network of four resistors and two capacitors is used to shift the dc level of the signal from +200 to -200 volts without distorting the pulse shape. A further example of our standard technology was the clock pulse gate. Our clock pulses were 1/3 microsecond wide every microsecond. These produced current pulses of the same form, with an average of 3 mA. They were capacitatively coupled from the anode load into the next cathode also with 3mA defined current, so producing pulses of 10 mA peak.

To return now to mercury delay lines. We started off with long and short horizontal delay lines, and subsequently we developed vertical delay lines. The initial long delay lines were all horizontal ones. Later on the folded long delay line was invented by Donald Davies. This was vertical with two reflectors at the bottom and the two crystals at the top. The signal went down one tube, across the bottom and up the other tube.

I next want to talk about the magnetic drum. We needed a backing store for the ACE Pilot model to provide us with more storage capacity and we wondered for a long time how best to do it. We said to ourselves "wouldn't it

be nice if every track on the drum stored exactly the number of bits in a long mercury delay line so that we could do a total transfer from one to the other in one revolution of the drum." We wished to avoid having a lot of buffer storage in the way, and concluded the only way to do this was to synchronise the drum to the machine. We concluded that if we ran the drum at one ninth of the speed of the mercury delay line we would indeed be able to transfer the whole contents one way or the other in exactly one revolution of the drum, being nine revolutions of the mercury delay line. We worked out a system where we could do this provided we could synchronise the drum to the machine with an accuracy of plus or minus four microseconds.

These are some details of the drum we finished up with: drum revolution: 9ms; bit time: 9 μ s; 1024 bits transferred in 9ms. It had 16 reading and 16 recording heads to access the tracks. They were shiftable to eight positions so we had 128 tracks altogether.

We subsequently went on to develop the big ACE drum. On this model the drum revolution time was 5 milliseconds, which meant 200 revolutions per second and the peripheral velocity is almost 200 miles per hour. The bit rate was 0.33MHz which is approaching the mercury delay line rate in EDSAC which was half a megacycle.

We subsequently went on to design the big ACE which was a three address machine whereas the Pilot model and the DEUCE were two address machines. It had two source highways and one destination highway and they were linked by a function unit. So you could take two words from two delay lines then, for instance, add them together, and send them to a delay line, all in one word time. Moreover we eliminated the time wasted between instructions so we could in the best instances carry out an operation every 32 microseconds.

The word length was 48 bits instead of 32, and this provided enough space to specify both sources, the function and the destination (as well as the timing information as used in the Pilot model). We had four drums on the ACE each with 256 tracks providing a capacity of 32,000 words of information.

Working party reports

Pegasus

John Cooper, Chairman

Since the last issue of Resurrection we have progressed to the point where our Pegasus is almost in working condition.

The alternator has been rewired by an external contractor, and has now been mounted on suitable blocks. It was tested and found to be working satisfactorily in June.

Since then, the air conditioning chiller unit has been installed outside the room containing the Pegasus. The computer itself has been assembled and sealed, to ensure that the circulating chilled air cools the machine and not the room! We have verified that this works too.

We have checked out the connections between the processor unit, the power bay and the mains distribution. We have also checked that the drum runs to speed, and that we have got the HT on the machine.

We have encountered one or two minor problems while getting this work done. There was non-standard wiring on the mains panel for the standby supply for the cooling fan - it had apparently been modified while the machine was at West Gorton. In the alarm system there were two thermal trips that did not, and still do not, operate. We had a vibration problem on the HT interlock, and two fuses blowing in Bay Two.

All of these problems are now cleared or being dealt with. The machine has been run, without the drum but under crystal control, and the monitors seem to be working. I anticipate that during our next working Sunday on 2 September we should be able to complete the commissioning.

We have made sufficiently good progress, then, for me to start thinking about the long term ramifications of running the machine. We need to have a policy on the alternator, for example, as this is a part of the machine that will wear out. We also found while investigating the fuse-blowing problem that some of the wiring in the top section of the computer is disintegrating. This is another problem that will get worse as the machine is used.

Elliott 803

John Sinclair, Chairman

Progress has been slower than we would like, principally because the Working Party members with knowledge of the 803 all live and work some distance from London. This effectively means that a trip to the Science Museum takes a day to make the effort worthwhile. The Working Party's biggest need is for workers experienced in the 803 who are based in or around London.

Work on restoring the computer hardware to working condition is still in the very early stages. We have started making an inventory of hardware components, and have also started on cleaning the edge connectors, with the help of Science Museum staff.

This is vital work which has to be done before we can start trying to make the computer work, otherwise poor contacts could create apparent fault conditions and make trouble-shooting very difficult. The 803 was in store for some time before the Society started restoration work, and most of the gold-plated contacts have turned black. There are some 60 on each board, so completing this task will take time.

We have also begun to sort out some of the books on the machine, mainly programmer manuals and user manuals. A start has been made on an inventory of paper tapes, of which a great many came with the machine. I am currently working on building a software emulator for the 803 to run on a PC.

We appear to have no shortage of parts. The Science Museum owns a second, incomplete 803 as well, and this should provide a supply of spares should we need them. The one problem is with the battery.

The 803 uses a specialised nickel cadmium battery, and will not run without it. The battery on our machine was naturally not maintained while in store, and is very badly corroded. In any case these batteries wear out with use, so we do need additional battery cells. The problem is that we do not know the manufacturer's name, so we have been unable to find replacements. Any help with this problem would be greatly appreciated.

S100 bus

Robin Shirley, Chairman

Activity is proceeding slowly at present. In some respects our activity is less pressing than that of the other working parties, as our area of interest is relatively recent and S100 systems are still in use, but it is important to get moving before people throw away manuals or software that is worth preserving. We have already received some equipment and associated manuals, mainly from John Cooper of the Pegasus Working Party to whom we are very grateful.

Two things are likely to happen before the next issue of Resurrection. First, we shall advertise for new members in an appropriate place such as one of the personal computer magazines. Second, we need to formulate a plan of action - what we are going to do, and what we need to preserve.

There were a fairly small number of S100 bus machines sold in the UK - mainly North Star Horizons, Cromemco Z2s and Vector Graphics machines. It would be nice to acquire one of the few original Altairs imported here, and also one of the Imsais. We are also interested in machines that were actually built, or at least substantially modified, in the UK.

Software

Martin Campbell-Kelly, Chairman

An evening open meeting for this newly-formed Working Party was held on Thursday 28 June. The subject of the meeting was "How to Tackle Historic Software and Data", and it was attended by some 40 members.

The meeting was opened by Tony Sale who explained the background to the formation of the Working Party. I then spoke about some of the main issues in preserving computer software and discussed the conclusions of a report on a "Symposium on the Preservation of Software" held in March 1990 at Columbia University, New York. The secretary of the Working Party, Chris Moller, then spoke on the preservation of data (as distinct from programs).

A lively discussion ensued, in which much attention was given to the value of preserving paper versus machine-readable records, and the need for emulators or working hardware to make software artefacts meaningful. The main conclusion was that the problem of preserving a significant proportion of software and data was entirely beyond the resources of the Computer Conservation Society.

It was, however, agreed that a "bottom up" approach of focussing on the software on a small number of early machines which are being actually

conserved by the Society could be usefully pursued. It was subsequently decided that one member from each of the machine Working Parties should be invited to join the software working party to help achieve this objective.

DEC

Adrian Johnstone, Chairman

In the Spring Tony Sale and I were interviewed for the DEC internal company newspaper which goes to all employees. As you might expect this brought several old DEC hands to our aid. In particular Richard Buxton, Mike Randall and John Willis (all PDP-8 aficionados) have attended a series of working party meetings at the Science Museum. The main focus of attention has been the Museum's original classic PDP-8. This was subjected to a very thorough checkover and was then successfully powered up. Several problems have been solved by swapping and repairing modules. Presently the machine will execute programs, but the core memory and associated sense amplifiers are not performing well. Many long hours have been spent tuning the memory, a process that will be completed by the end of August. We now have a good supply of spares after salvaging a system discarded by Harwell.

We have also acquired several other machines, mainly through the good offices of John Willis. We have a working WS78, a DECsystem 300 and a 310. These are large systems (well, large by minicomputer standards anyway) with hard disk drives. The 310 is running but with one faulty drive.

Also under the DEC working party wing Max van Daalen has donated and restored a Philips P352 office machine and Peter Hoare has donated an Olivetti P652 system. Hatfield Polytechnic, the Polytechnic of Central London and London University have donated equipment including teleprinters, PDP-11 and PDP-8 parts and a quantity of Digico minicomputer equipment.

Our next large project is the restoration of the Museum's PDP-12 which is a hybrid PDP-8/LINC machine designed for applications requiring large scale analogue interfacing.

Letters to the editor

Dear Sir,

I would like to say how delighted I was to hear of the formation of the Computer Conservation Society and how interesting I found its first bulletin.

I was, however, disappointed with the omission of LEO Computers from the text of Dr Campbell-Kelly's talk entitled "ICL and the British Computer Industry"

LEO Computers made a tremendous contribution to the British computer industry and had a significant influence on the development of ICL. LEO produced the first commercial computer and, with LEO III, the first true multiprogramming commercial machine. A number of LEO innovations were to appear in the 2900 series and indeed even today 30% of this Society's mailing list are still employed by ICL.

I would like to reiterate my full support for the CCS's work, despite my disappointment about this article, and will try to attend future meetings.

I am, Yours sincerely,

Peter Byford, MBCS, Chairman, LEO Computers Reunion Society, Ware, Herts

Dear Sir,

Having just received a copy of the first bulletin of the Computer Conservation Society, may I thank you and all the contributors for such an interesting and readable journal. It is good to have such an open and honest publication with reliable reference data.

If there is a cost for future issues, please let me know. Also can you please advise me of the cost of further copies.

I very much look forward to another excellent issue. Yours sincerely,

D Len Peach, Human Factors & Historical Archives, IBM UK Laboratories Ltd Hursley Park, Winchester

Editor's note: Copies to non-members and additional copies for members are available at a price of £3.00 each, or £10.00 for a complete year's set of four issues. Interested readers should contact Tony Sale at the Science Museum.

Skills and Objects Wanted

SAFT Nickel Cadmium battery for the Elliott 803
Paper tape copy of Tony Hoare's ALGOL compiler for the 803
Any Williams Tube memory bits, particularly from the IBM 701, 2 or 4
Enthusiasts for the Totalisator Working Party
Experienced Elliott engineers and users

Forthcoming Events

Thursday 11th October

CCS Open Day at the Science Museum, sponsored by *Computer Weekly*.
10am to 5pm.

An opportunity for you to see for yourselves the progress made on restoring the Society's old computers, and hear presentations from each of the Working Parties. Non-members are welcome, and membership forms will be available. The meeting point is the Lecture Theatre.

Thursday 15th November

All day meeting in the Lecture Theatre of the Science Museum, celebrating the thirtieth birthday of Algol. Starts at 9.30am.

1991

A programme of meetings is currently being planned by the Committee. These will take place on the last Thursday of each of January, February, April, May and June 1991. More details in the next issue.

Resurrection is the quarterly 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

Production Editor and typesetting - Adrian Johnstone

Cover design - Tony Sale

Printed by the British Computer Society

Committee of the Society

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

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

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

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: 071-477 8000

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

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

Chairman, S100 bus Working Party **Robin Shirley**, 41 Guildford Park Avenue, Guildford, Surrey GU2 5NL. Tel: 0483 300800

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

Editor, Resurrection **Nicholas Enticknap**, 45 Spencer Court, Spencer Road, Wimbledon, London, SW20 0QW. Tel: 081-947 2812

Anne Berne, 16 Nevem Square, London SW5. Tel: 0483 300966

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

Dan Hayton, 31 High Street, Farnborough Village, Orpington, Kent BR6 7BQ. Tel: 0689 52186

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