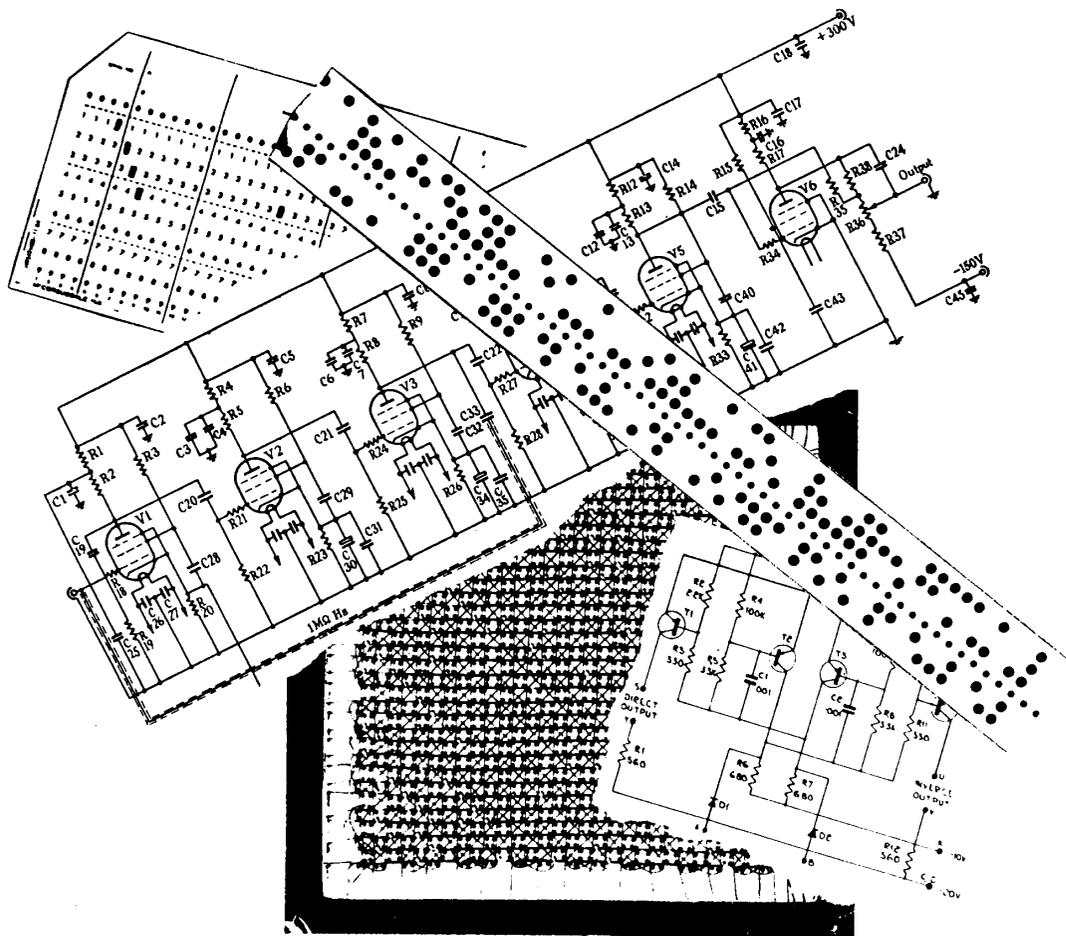Computer

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.

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

Resurrection

The Bulletin of the Computer Conservation Society

Number 10

Summer 1994

Contents

Society News
Tony Sale, Secretary

Things have been progressing well for the Society at Bletchley Park. In contrast, little progress has been made at the Science Museum after the move to Blythe House.

The meetings programme has suffered a number of cancellations, primarily because we still do not have a meetings secretary, and I have been too busy with Bletchley Park to organise meetings at the Science Museum.

Bletchley Park does look as if it is going to come good for the Society. The Bletchley Park Trust now has a two year lease from British Telecom on H block and Faulkner House. This has allowed the computer exhibition and the Colossus rebuild project to really move forward.

Minds were concentrated by fixing the Opening Day for 18 July 1994, with the Duke of Kent doing the honours.

The Elliott 803 has been moved from Andrew St Johnston's barn into the computer exhibition room in H Block. After being stripped down it is now showing signs of life after more than 12 years inactivity.

A Burroughs L5000 visible record computer has been acquired from Sir Robert McAlpines and has been fully renovated by Unisys, thanks to the company's PR director Martin Sexton. This computer was bought new by McAlpines in 1974 and was last used in January this year, a very impressive record.

We also have an IBM 1130 which needs restoring to working condition.

An archive room has been allocated in Faulkner House. Archiving work is set to start there soon after the Opening Day.

The next open weekends in Bletchley Park are 6-7 August, 20-21 August and 3-4 September. We need a roster of CCS members to man the Computing Exhibition room on these days. Please contact me if you can help.

New contact point

Readers wishing to contact the Secretary are reminded that he is now running the secretariat from home. The secretarial telephone number is 0234 822788. Letters should be addressed to 15 Northampton Road, Bromham, Beds MK43 8QB.

Evolution of the Ace drum system
Fred Osborne

I joined the Ace team at the end of 1948, within a few months of its formation from the merger of the maths division and the electronics section. I was charged with looking after what my electronics friends choose to call the ironmongery. I'd like to mention by way of introduction a couple of other pieces of Ace ironmongery which were evolved before the drums.

When I arrived many of the electronic units of the Ace had been tested on the bench, including the circulation unit for the mercury delay lines, but none had been put together as a working system. Once this was done, problems with the delay lines became apparent.

The original delay lines were horizontal tubes with the crystals mounted on flanges, one each end. They stood out in the open, so changes in temperature caused changes in length and variations in delay with consequent problems. This was overcome quite simply by putting them into a wooden box with a thermostatic heater. However, that brought further problems as the crystals had a habit of dying. To fish the tank out of its box, tip out all the mercury, refix the crystal, and put it all back again was a major upheaval.

So the search was on for a better, more compact system, and several ideas were tried. One was a thin square box on its edge filled with mercury, with the heads sticking out at the top like horns so they were immediately accessible for changing crystals, and mirrors round the inside of the box to give you multiple reflections. The box was about 15" square and 2" thick. The idea was you could stack them together like books on a shelf. It made a nice compact unit and each one could be serviced easily. But this brought its own problems and other solutions were proposed.

The eventual winner had the line divided into two parallel vertical tubes, with a reflector system at the bottom. The crystal heads were at the tops of the respective tubes, in a mounting system designed by Donald Davies. This unit provided for adjustment of the length of the line, and also for rotation in two planes to line up the crystals. This was the design used on the Pilot Model, the Deuce and the final Ace.

The other piece of ironmongery was the Ace Pilot Model itself: to construct it we needed to put all the units together in a sanitary sort of fashion. With all the problems with the mercury delay lines we became disenchanted with the idea of building the 1024 of them required for the final Ace, so thoughts turned to other recording media.

After one or two exotic ideas had been explored and discarded, magnetic recording was the winner. Then we had to decide: did we want endless loops of tape, or discs, or drums? The drum won that battle. If we were to make a drum which was truly a replacement for 1024 delay lines it would have to have 1024 tracks, and one revolution per milli- second. That meant 50,000 rpm, 1024 reading heads and 1024 writing heads. It would have been a fearsome thing to embark upon.

The eventual solution was to have a lower capacity drum, synchronised to the computer so that further drums could be added to build up the total data storage required. This decision ruled out synchronising the computer to the drum, which otherwise would have been easier. We thought about 5000 to 6000 rpm would be a reasonable speed to aim at, and perhaps 32 or 64 tracks might be on. In the end we managed 256 tracks on the final Ace drum: whether we achieved 256 tracks on the Pilot Model I can't remember.

So to the design of the Ace Pilot Model drum. The Pilot Model at that time was conceived as a test bed for the final machine, and the design of the drum followed this philosophy in that any component - the motor, the drum itself, the head shifting mechanism - could be changed independently without interfering with the rest of the subsystem. It was not an integrated design where if you move one thing everything else had to be moved as well.

We chose a speed such that one drum rotation occupied exactly nine circulations of a delay line, which worked out to a speed of 6510 rpm. The individual binary digits were written to the drum at 9 microsecond intervals, so the 1024 digits in the delay line exactly occupied one track. Picking digits at this interval meant that the digit 1 from the delay line was followed on the drum by 10, then 19, and so on. Since the numbers 9 and 1024 are mutually prime, all the digits were eventually inscribed once each, though in a scrambled sequence. The reading mechanism, using the same picking arrangement, automatically returned the digits to the delay line in their original natural sequence.

One of the problems with the magnetics was to make a reading and writing head occupy the same space on the drum. You can do it by switching the head, but when you bear in mind the enormous difference between the reading and writing current it presents some nasty problems. These were avoided by observing that in the scrambled sequence mentioned above, digit 2 from the delay line was separated from 1 by about three-eighths of the drum circumference: the writing head was mounted at this displaced position, and the gating signals delayed by one microsecond so as to communicate with the correct digit position in the delay line.

To run the drum in synchronism with the computer required a synchronous motor with a good strong synchronising torque. Looking through the literature we discovered a thing called a hysteresis motor which is a delightfully simple device. The stator is a conventional three phase sine-distributed winding, but the rotor is nothing more than a cylinder of Alnico permanent magnet material. It runs up as an induction motor but as it gets to synchronous speed the flux from the rotating field builds up a permanent salient pole, as it were, in the magnet and so there is a very strong synchronous torque.

The normal way of making these magnets was to cast them. Unless you X-rayed every one you couldn't be sure there were no blow holes there, and I can remember the first one just shattered, causing a bit of havoc with the inside of the stator. But we managed to get some forged and these were very successful.

The motor was hung on the bottom of the drum itself, so the motor could be changed.

We could have put one head per track and switched heads but this seemed an unnecessary complication. We opted very early on to have a small number of heads and to shift the heads parallel to the axis of the drum, so that one head could write different tracks. We had 16 heads and we could move them to 16 positions, giving 256 tracks.

The next problem was the head shifting mechanism itself. After several ideas had been discarded one due to David Clayden was tried. The output to the head had three solenoids coupled together one above the other. When the first solenoid was energised a gap closed, causing a plunger to move up on the stock a distance of one track width. If the next solenoid was energised the second gap was closed, a distance of two track widths. If both solenoids were energised the plunger moved three track widths. The third solenoid moved its plunger four track widths, so that we had eight possible combinations and tracks. We tried to get 16 but that pushed the technique too far and it wasn't successful.

To hear this thing working was quite amusing because it would clatter like a lot monkeys with false teeth on a cold night! And you got some rather interesting rhythms. This drum was used on the Pilot Model during its time as a working computer and was quite satisfactory.

Drums of this period invariably carried a lot of scratch marks. This was quite a problem. As we got more and more adventurous and screwed the heads closer to the drum to get a little more signal, eventually they touched,

and when they did they scored big grooves. This led us to look at a way of producing a tougher coating.

The solution was to mix the powder with an epoxy resin, spray it on with an artist's spray gun, then let the thing set. This produced a tougher coating which resisted minor crashes, and if it was a really bad crash you could repair it by painting the coating on with a brush.

As we got better and the heads became ever closer other problems arose. One was the unevenness of the coating because it was sprayed on by hand with an artist's brush. The answer to that was to mount the drum in its bearings, take it into the workshop and put it on a milling machine, put a diamond tipped cutter in the place where the milling tool would normally be, and then run the drum at a relative- ly slow speed past this diamond tool so it skimmed a slight skim off the surface of the drum to establish a datum. Then you sprayed on the coating as before, and after it was cured set the tool back by the required amount and skimmed it again, so the surface of the coating left was a known thick- ness and was uniform. This was very successful.

As always, you cure one problem to expose another. The bearings we were using were a commercial commonly available type, but the run out of these bearings was now of the same order as the head gap. Fortunately we found a firm which had grown up during the war by developing a technique of re- working commercial bearings to the higher tolerance needed for gyroscope work. They provided us with some doctored bearings which we used with great success.

Then we began to think about bearings and their problems and designed one of our own. We proposed it to this firm: they made us some samples which they liked and put into their catalogue. So the Ace drum had bearings of our own design.

Later we used another type of head shifter. It was really a moving coil analogue of the solenoids. It contained a conventional loudspeaker magnet but with a gap at each end.

In each gap there were two concentric coils. One coil moved between the top of the pole piece and fixed stops, giving you a two position shift; the other coil moved between the underside of the original coil and the stops carried by the first coil, analogous to armatures in the solenoid. There was an identical arrangement at both ends. That gave us positions 2 and 8, and by the 2:1 reduction 1 and 4, so you got the 1, 2, 4, 8 shifts. It was marginally quieter than the original and was used on the first Ace drum.

By this time the Deuce had been installed and up and working. The Deuce drum was similar to ours except it had a servo controlled head shifter with 16 tracks, engineered in a very elegant and satisfactory way. So we abandoned the digital head shifters, and the final Ace drum used a servo controlled shifter. The Deuce drum used a potentiometer for position sensing but we used two synchros in a two term servo control of the head shift. That's how all the final Ace drums were made.

In conclusion, the final Ace drum was run at a much higher speed than the original. In fact it was 12,000 rpm which, bearing in mind the 1536 digits in a delay line rather than 1024, enabled us to run with a five sector drum instead of a nine sector as on the Pilot Model, but essentially it was the same arrangement.

On the Pilot Model we had separate head shifters for reading and writing but programming experience indicated that this wasn't necessary. This eased the problem of having to get the two movements to be exactly the same because they were coupled together. On the Ace drum they were coupled by tapes which moved the heads up and down.

One of the problems I've mentioned was the scraping of the drums because we had no real way of knowing how far away the heads were from the drum surface. The gap was only of the order of 0.4 thou which is rather difficult to adjust accurately.

Friends in the workshop who knew about small measurements used to provide us with packets of cigarette papers to use as feeler gauges - a toolmaker's traditional way of setting these things up. On the final drum we had an air gauging system, with a jet at top and bottom of the stack.

The air gauging system was rather like a pneumatic analogue of a Wheatstone bridge: you compare the gap between the drum and the heads with the pressure developed across a fixed orifice by calibrating the manometer with a micrometer beforehand. With it we were able to set the things up quite easily to the order of 0.4 thou without touching the drum. In the end we had four drums running synchronously with the machine.

I forgot to mention the toothed wheel. The wheel had one tooth for every digit position on the drum. These produced what we called drum pulses at one-ninth of the clock frequency of the machine. The drum drive was generated by a ring of three oscillators whose average frequency was locked to one-ninth of the clock speed, but we needed to be sure that the position of each digit on the drum was known. This was done by comparing the phase of the signal derived from the toothed wheel with the counted down signals from the clock, and injecting a phase advance or phase retard into the drive

amplifier. Each digit was monitored and corrected in that way. The drum pulses were also used elsewhere in the logic for transferring information to and from the dedicated delay line.

This article is based on a paper presented at the Society's all day NPL/English Electric seminar held at the Science Museum on 20 May 1993. The editor acknowledges with gratitude the assistance of George Davis in editing the transcript.

Pegasus simulator

Members of the Society who have early versions of the Pegasus simulator PEGEM may be interested to know that the latest version has been much refined, with improved graphics and an online manual for operation of the simulator. If you would like a copy, please send four 25p stamps to cover the cost of the diskette and postage to Chris Burton (for address see inside back cover). The simulator runs on any industry-standard PC with VGA graphics and a hard disc, and is supplied on 3.5" diskette unless 5.25" is requested. In the latter case, please state whether 360Kb or 1.2Mb format.

Memories of the Manchester Mark 1
Frank Sumner

> This article is an anecdotal reminiscence of the author's experiences with the Manchester University Mark I computer.

In 1951 I started a PhD course in Computational Chemistry. Sometime after Christmas I was doing calculations on energy levels involving hydrocarbons.

I was lucky, I was using a Marchant. There was a motor in the back, so you didn't get tired, you just got bored. My supervisor said "why don't you go and see Alan Turing - I think he has got some sort of computer over the other side of the campus".

I duly went, and said that Christopher sent me over to discuss using the computer. Turing's response was "Here, read that", handing me a book. It was the Mark 1 Programming Manual, which he had written himself. "Then go and see Cicely Poppelwell, who will give you some time on the machine, and start programming." End of conversation.

I disappeared for a couple of months, and discovered that the perfect way to write a tutorial textbook is to fill it full of examples, none of which work. That might be an exaggeration - let's say most didn't work. Turing was like that: if he wrote 'k' instead of 't', he knew it was 'k', so why bother to do all that proof reading? So all the programs written in Mark 1 code had slight errors in them, and by the time you had worked out what the code should have been you had become quite a competent programmer.

You then constructed your program. You know all the things you are taught to do nowadays: write a small module with clean interfaces, employing top down design; test your procedure; embed it; test it some more - all the good stuff we still try to teach students. I didn't do any of that. They gave me a problem and I simply wrote a program to print out the answers and took it to the machine. A few months later, it worked!

To achieve that I acquired a copy of the second edition of the manual, which was written by Tony Brooker. It is the most glorious piece of understatement you could imagine. At the beginning it says, "Much material has been taken over and altered, or only slightly modified, from the first edition written by AM Turing". `Slightly modified' meant about one character in every 20.

Brooker acknowledges Cicely Poppelwell, Nick Hoskin, Alec Glennie - names some readers might remember. There wasn't a big group working at that time.

Occasionally people claim to be able to talk hex. Well, that involves just four binary digits. The Mark 1 had 40-bit words organised into 20-bit lines, so there were two lines per word. These 40 bits were arranged in groups of five, with one character to represent each group of five. The user was strongly recommended to memorise each of these characters.

The real reason we had to learn them was to punch certain odd characters on to paper tape. These characters weren't printed when they were punched in, so if anyone interrupted you halfway through, you had to find out where you'd got to by identifying these characters on the paper tape itself.

Another nice bit comes at the end, where you are strongly recommended to learn these characters because when you are putting in decimal numbers, you can't type them in decimal: it is much quicker to enter the numbers in binary and convert them to decimal in the machine. The next line says "This can be done quickly with a Brunsviga".

That is a really nice line for some one who was going to start using a computer which is massively faster than a Brunsviga. Remember that the speed of multiplication on the Mark I, compared to anything previous, was a step function greater than that achieved by any subsequent machine.

On page 9 of the Introduction the manual starts telling you how numbers are represented. The first example is 10 which tells us nothing because it is symmetrical. The next one is 26, and you think "How Come?". Because up to now the book hasn't told you where the least significant digit is, and which way significance changes. I gather that Turing gave lectures and would do normal mathematics with decimal notation `back to front' because he had got used to it. But perhaps that is an apocryphal story.

When these days you hear about risc machines, you are usually told this stands for *reduced* instruction set computers. I regard them as *rediscovered* instruction set computers. You can't get much more reduced than the instruction set we had.

There was no divider: that was done by software. The multiplier, designed by Alec Robinson, took half of the valves in the machine. There were so many valves in it that nobody thought they would keep going for long enough for a multiplication, though in fact they obviously did.

What we did have was a most significant digit instruction and an instruction that counts the ones in a word. These were very useful to some users - the CDC 6600, the mid sixties supercomputer, also had these instructions.

Turing also got a random number generator built in. The numbers were really random, not generated by a repeatable sequence of operations. If you had a case which you wanted repeating, the second run would create a different set of random numbers!

One of *the* great inventions of the machine (apart from it being a memory or stored program computer with a massive memory of a few thousand characters) was address modificat- ion, which in those days was called `B' modification.

The console had two displays which you could switch to any of the eight storage tubes. You could look at the bit pattern - a lot of dull and bright dots - and read them. For example, four bright dots and a dull dot was a `K'.

The modifier was called the `B' tube. There was a divider which contained the two numbers being divided called the `D' tube, a control tube `C' and an accumulator `A'. So when they built these index registers on another little tube it was natural to call it a *B tube*. The contents were called *B lines* and the process was called *B modification*.

I can just about stop calling it that to the students today. They look a bit askance when you start talking about *B modification*: they think you are using bad language.

For modification three of the bits in the instruction indicated the B-line to select. The instruction had 10 address bits and one would have expected that the B-lines would also have had 10 bits to provide address modification. However the B-line in fact had 20 bits, the same as the total instruction, which allowed altering the function field as well as modifying the address field. This design feature permitted some rather exotic instruction sequences, whose effect bore little resemblance to the code written down. Luckily there was no further modification of the modified instruction!

One of Tony Brooker's contributions to compact programming was the routine changing sequence (essentially the operating system), which was only a few lines long.

One of our less sensible moves was storing the powers of two You've got a tiny memory, you squeeze everything into as few lines as possible, and what do we waste 10% of memory on? The powers of two! It's nice to know that

most people make mistakes: every time you mess something up now just look back into history and you will probably find a similar mistake.

The tubes were regenerating all the time by rescanning, and the digits were bright for ones and dull for noughts. When you read a line you scanned it at the same time as the re- generation scan. If you read it repeatedly that line would glow, getting progressively brighter. So when you got stuck in a loop it wasn't difficult to find it; you just looked at the tubes till you found a bright band. You could watch these bright bands moving round the program and you could say "another 10 minutes to get to the last stage of the program".

The cathode ray tubes (CRTs) were particularly sensitive. The trams that went past the Dental Hospital sometimes produced sparks that caused interference with the machine, changing 0s into 1s.

The digits in the top right hand corner of the CRTs would start changing from 0s to 1s, and if you got an empty store you could actually see the corner getting whiter as the digits crept in. So you got to the stage where you didn't put numbers and constants near the top right hand corner of your page. You could sit watching and say you had another few minutes before they got wiped out.

It was all fixed point arithmetic obviously. Who wanted floating point? When they invented it in hardware for the Mercury an awful lot of people said "you can't do that, you won't know how accurate your answers are". Luckily the engineers said "we're not bothered how accurate the answers are, we just want them quickly". But the mathematicians really objected to the fact that you didn't know the accuracy of the arithmetic.

One beautiful drawback was: if you added two numbers together using twos-complement notation, you could add two positive numbers together and get an overflow into the sign digit and get a negative answer. That was your fault - the machine didn't bother to tell you, as it did not detect overflows.

The manual doesn't refer to the magnetic drum, but to the *magnetic wheel*. That gives you some idea what shape it was. It also refers to digits being represented as areas of magnetised nickel. Because ferric oxide was a bit cheap and nasty, we got a beautifully nickel plated drum, which was good in a sense: when the heads crashed into it you didn't gouge a big hole, you just put a tiny scratch on and you could still use it.

At the entrance to the machine room there was a notice listing all the disc tracks: against each track it would say good, bad or dodgy. That table was changed daily! So on the front of your program was a translator to change the

tracks in the program to the tracks that were good that day. Dave Howarth has pointed out that this was the first use of virtual memory.

If things went wrong it could be physically dangerous. A piece of metal could fly off the fast printer, so an old army tin helmet was available to put on before going into the room containing the printer.

You got used to working long hours, generally from 8am to 8pm. Of course you had to keep stopping when the computer went wrong and the engineers tried to fix it. At the end of a night's work it was questionable who had done the most work - you or the engineer. Many close friendships were formed between users and engineers: some even led to marriage!

One of the hazards at the end of a night's good computing was Mrs Dixon. She would occasionally push her polishing machine into the power supply, and one night I went out to find her in the corridor saying "Have I broken it?".

Ferranti's first programmers were all very attractive young ladies recruited to the group by Dr Prinz. Many of those young ladies married members of the programming team or engineers. Tony Brooker's wife Vera was one of those first programmers.

This article is an edited version of the talk given by Professor Sumner to the North West Group at The Museum of Science and Industry in Manchester on 21 October 1993.

Editorial fax number

Readers wishing to contact the Editor may now do so by fax, on 081-715 0484.

Ferranti in the 1950s
Charlie Portman

I joined the Ferranti Computer Department in 1954. I joined as a Circuit Engineer: there wasn't anything like Computer Science at university in those days. There were people in Manchester and Cambridge working with the Manchester Mark I and Edsac, but there were no Computer Science courses, and I only became interested in computers by accident. I was doing an Electronics Course at Liverpool and I happened to go past a lab where work was being done on an Analogue Multiplier: that was what got me going.

So I became a Circuit Engineer. It was quite a long time before it was generally recognised that there was a another task, that of Logic Designer. The number of job functions has escalated several levels since, and there is now a whole range of different engineering skills involved in building machines.

I was initially involved in test programming. Then Peter Hall transferred me into the works, into the drum commissioning lab. Although Ferranti had designed and built these wonderful drums, the works couldn't actually make them, so I got involved in that.

Later I worked on Sirius, Ferranti's first attempt at a production semiconductor-based computer. Then I worked on Orion, then on the 1900, and in 1972 I changed from a Hardware to a Software Engineer.

In trying to put a structure to all this there are three things that I would like to talk about and to emphasise.

Firstly the structure of the design teams was not fixed. In a sense we were not only developing computers, we were inventing computing as a profession. A lot of computer engineering practice we developed ourselves, along with other teams in other companies working on similar projects.

For instance when we built Sirius we took a piece of paper about the size of a table and we drew the whole machine on it. We had to write fairly small to do it, but we got it on one piece of paper. The concept of levels of abstraction or dealing in units didn't exist then.

When we came to build Orion that approach proved to be impossible. In fact to my knowledge there was only one man who could carry the whole of Orion around in his head, and that was Ray Levitt. He knew the names of every package in the machine, what it did, and what the wave forms should look like.

So there we were, coming from our various backgrounds. We were Circuit Engineers, or Mathematicians, or Chemical Engineers. We were coming together to create from the sort of thing the universities had pioneered, but with a lot of invention of our own: trying to build machines that one day would enable the company to make some money.

So that is one thing I would like you to remember about the fifties: it was then we invented the concepts of the organisation of design teams, documentation, interfaces, requirements specifications, all the good things which have gradually developed, perhaps taken from other types of engineering, into the Computer Design Process.

When we did circuit design we used slide rules, or if it was really difficult log tables. We didn't have calculators. When we tried to design big computers and tried to keep track of all the logic elements and the interconnections, we didn't have computers to help us to do it. It all had to be done by hand, by draughtsmen in the drawing office quite often. In fact I would like to lay claim to be one of the inventors of the concept of Design Automation.

I was responsible for Sirius and we managed to get the wiring and the positioning of the packages and the drawings all out of step. So Gordon Calverley and I hung one of the machines on a crane in the works. We sat behind it and wrote down every wire, then we went round to the front and wrote every package. Then we went to a certain senior programmer and asked him "Can you program a computer to do this?". He did, and the Pegasus at Hollinwood thereafter ran a piece of software called Wirelog which tried to keep track of those things. I don't know how we would have managed without it. The task had just got too big, it was out of hand.

So that's the second thing: we were inventing not only organisation techniques and computer engineering skills but also various ways of using the technology we were developing to help us keep track of the enormous numbers of components and interconnection and so on.

The third thing I would like to say is that is was fun! We worked all sorts of hours. I think when I started we worked Saturday mornings: we were on a five and a half day week. I don't know how many hours that was. But we really enjoyed doing it and did all sorts of unpaid overtime.

Some of the things that have changed a great deal are to do with the change from valves and high voltages down to transistor level. You can still get quite exciting effects from the power supplies on large transistor machines, but they pale into insignificance compared with taking a pair of

wire cutters and putting them on top of a Mercury, which had three copper bus bars with 400 volts 50 cycle AC on them.

I can remember a Canadian gentleman who was a bit of a japester. When we first got semiconductor diodes one of his favourite tricks was to paint them with repeated coats of transparent glue till they formed a big blob. Then one night when he was working late and you weren't, he would wire one of these backwards into your rack of equipment between 300 volts and earth. When you came in next morning and switched on there was a spectacular explosion because the plastic shell round the diode held the gases in until it couldn't stand it any more.

The way drum head clearance was judged was by turning the screwdriver, whose end was held in your ear, carefully, until you could just hear the head just touching the drum then back it off a bit. The nickel used was strong enough to do that, if you were careful. There is a famous cartoon with a picture of a drum with a slice cut out of the top and someone asking "Is it touching yet?"!

When we switched over from nickel to oxide, which we did because we thought we would get much bigger signals from it and so be able to produce a more reliable storage medium, one of the first things we discovered was that the oxide would not stick on.

You took this drum which you had lovingly engineered to the paint shop to get it sprayed with this brown gunk, then you checked it with a micrometer and found it was full of bumps. So we tried to devise a method of turning it with diamond tipped tools to get it flat. Sometimes it worked but quite often the whole of the surface would come off in one sheet and float gracefully across the shop. So we developed all sorts of techniques for etching the surface so that the oxide coating would grip.

If you got a head crash on an oxide drum it was almost always fatal. The Manchester Pegasus shows interesting signs of that: luckily, we haven't lost the whole surface but there are a number of badly damaged tracks.

One other thing about the drums: we were obsessed with getting the weight down. They were quite heavy, originally brass or some proper traditional British material. We wanted to use aluminium, but it wasn't strong enough, so we went over to a magnesium alloy, and this had some interesting effects. The first time we tried to turn one it burst into flames, which was quite exciting.

This whole business required getting involved in Mechanical Engineering: devising ways of producing very small clearances; dealing with

thermal effects such as heat expanding the drum and doing away with the clearances; finding ways of measuring the distances that didn't involve you in contact (we developed some weird and wonderful techniques for doing that).

Another issue relating to drums was the access time. These drums went round at a very ponderous speed of 35 milliseconds and the programmers were always grumbling about how long it took to get the next page off the drum, so we tried very hard to get the speed up.

In, I was going to say West Gorton but we were probably still in Moston, we were very sure we didn't want motors with brushes on them. With our experience of the trams outside, the idea of having sparking brushes inside the drum was not at all appealing. So we ran with induction motors, and if you wanted to go faster than 3000 rpm you had to have a higher frequency.

I was given the task of pursuing some 1920s Italian Railway research into arc lamps to see if I could generate 150 cycles from a static converter, relying on overloading magnetic circuits and generating third harmonic distortions and so on. We accidentally generated what turned out to be a current source: can you imagine a three phase power supply driving a thing that generates two amps regardless of the voltage? This was quite exciting.

We used ordinary 100 watt electric light bulbs as ballast loads, and it generated two amps into the cold bulbs which heated up very rapidly. The voltage ran straight up to 1500 volts at which point the bulbs all went pop, and then there was a nasty brown smell. We did get it tamed in the end.

PCs do occasionally go up in flames, but very seldom. When they do, there is usually a little bit of scorch, but you don't get this sort of excitement any more.

What we have probably still not learnt to this day is communication across disciplines. I remember the mathematicians wanted to know how many bits at one end of the word were the same as each other. We kept asking why they wanted this but all they said was it was very important. So we built this facility in. It turned out they wanted floating point, and what they wanted to do was to standardise the number. Had we known this we could have built what they wanted more easily. This lack of communication applied in a lot of areas.

Another development which was taking place at this time was the addition of peripheral equipment. When I started, Ferranti was building the third Mark 1 star. The input medium at that time was five-hole paper tape punched on teleprinters, with the output emerging also on paper tape and being printed on the teleprinter. Other companies were using punched cards.

During the period we were building the later Mercury, Pegasus and Perseus systems we were starting to get magnetic tape, great big line printers, and card punches of course.

I recall Preston Council buying a 1900 to replace its ICT Tabulator. We installed the machine and the programmers did their bit and set the machine going. It punched a card, then read it and printed a line, and everybody was delighted, because that was what the previous machine had done. So, all three ran in beautiful synchronism. Unfortunately, the printer which ran at 1350 lines a minute then had to wait. The reader which read at 300 cards a minute also had to wait. The punch only punched at 100 cards a minute! So much for a machine with autonomous transfers, we thought!

I could tell you many stories. I could take you through the saga of trying to find the short in the back wiring in the Turitz Orion, that shorted the one minute interrupt in the clock onto the peripheral address parity bit: it took us ages to find that! I could tell you about John Allenby's wonderful gadget for connecting to the Ericsson keyboard, which used to go off like a battery of Howitzers periodically and pepper the room with transistor cases.

I could tell you about the time we were sitting exhausted in the Turitz canteen. I was eating my 30,000th cheese omelette (I am a vegetarian, which isn't a good thing to be in Sweden), when an Italian voice came over the loud speaker saying "The magnetic tape is working". All the Ferranti people immediately leapt up, dashed out of the dining room and down five flights of stairs.

It was a lot of fun!

Charlie Portman is Chairman of the North West Group Pegasus Working Party. This article is an edited version of the talk given by him to the North West Group at The Museum of Science and Industry in Manchester on 21 October 1993.

Very early computer music
Donald Davies

Most readers will be familiar with the National Physical Laboratory's Pilot Ace and full Ace computers for their contribution to the development both of computer technology and of the UK computer industry, as discussed in detail at the Society's all day seminar in May 1993. In this article Donald Davies describes one of these computers' less well known features, their ability not only to play but to compose music.

The Ace Pilot Model and its successor, the Ace proper, were both capable of composing their own music and playing it on a little speaker built into the control desk. I say *composing* because no human had any intentional part in choosing the notes. The music was very interesting, though atonal, and began by playing rising arpeggios: these gradually became more complex and faster, like a developing fugue. They dissolved into coloured noise as the complexity went beyond human understanding.

The origin of this phenomenon was a set of diagnostic devices which had been put into the machine by myself and David Clayden. We didn't ask anyone if they should be there because they cost so little and might be useful.

A cardinal feature of the machines, due to Alan Turing, was *optimum programming*. This was designed to make the utmost use of the (problematic) access to data and instructions coming from the long delay lines at specified times. If you did well, transfers of data would be taking place most of the time.

So a measure of programming quality would be the proportion of time that transfers were going on. I thought to myself that this could be measured very easily by smoothing a trigger called Transtim (a word, like its partner Timci, which is burned into the brains of those who knew the Ace). I had to remove the AC component from the signal to stop the needle flickering too much, and it seemed logical to put this AC part into a speaker just to hear what it sounded like.

The small meter on the control panel was a very crude measure of a very complex activity and hardly useful, but the sound source naturally led to amateur compositions and the customary transcripts of JS Bach.

The sound had a practical side. We were doing mathematical work which could be speeded up by suitable intervention. This might mean knowing when to stop or change the process, and sometimes the sound would tell the mathematician which procedure was in operation and help him decide what to

do. Remember that there was no display to write messages on, just punched card output.

When the Ace Pilot model was due to move to a new location and become a service, we thought that the crude old control panel should be replaced by a nice new one. I designed the working part and got a cabinet maker to construct the wooden case and desk.

David Clayden suggested some new switches that might help in machine testing and program debugging. These would enable any field of the current instruction to be changed. This facility would be used when the machine was working "one-shot", step by step. The number set on the switches for a particular field could be substituted for the one just received from memory. But if the machine were running full tilt, all the instructions would have the chosen fields fixed at their substitute values - what a mess this could make!

As soon as this new facility was working David characteristically began to experiment with misusing it, and this is when the music came forth.

I cannot remember the details but it was something like this. A long delay line with 32 instructions was the source of the program and probably it was initially empty. The program would be in a loop and therefore give a note from the speaker. By fixing the operation to add into a suitable bit in the instruction delay line, this counter would spill over into the timing fields and periodically enter a new kind of loop. Using the right switch settings made the loop size stay constant for a while (a note), then abruptly change in size.

Loops were always multiples of 32 microseconds long, so notes had frequencies which were submultiples of 31.25 KHz. The music was based on a very strange scale, which was nothing like equal tempered or harmonic, but was quite pleasant.

Initially the playing was rather slow, as though it was simply familiarising the listener with its unorthodox scale. Higher up the scale, when the notes are a bit high even for young ears, the spilling over process gets more complex. It is next to impossible to work out what is going on in the delay line, with the instructions changing all the time. The next bars of music are a little faster and the same scale of notes appears, generally rising but now varied in sequence like an arpeggio. This repeats with ever faster and more complex patterns until it becomes too fast and complex to follow. Once you get used to the strange scale it really sounds musical and even intelligent.

We were amazed when we first heard it. To get a nice composition you had to choose the right setting on David's switches, and the settings for the

interesting music were rather critical. David has reminded me that some switch settings gave falling arpeggios.

The Ace proper had a different instruction format. We built in the same diagnostic features and played the same music- composing games but somehow I think the Pilot model was the better composer.

Obituary: John Gray
Doron Swade

It is with great sadness that we learned of John Gray's passing on 28 April. John Gray first made contact with the Science Museum in the mid-1980s when his working Elliott 803, housed in a workshop attached to his house, could no longer justify the space it occupied. I first met him when I visited the 803, and subsequently spent an instructive Sunday crawling around on the floor under John's informed and gentle direction, uncabling the machine, labelling the terminations for subsequent reassembly, and sifting through the cans of magnetic tapes and 803 spares in his garage. It is John Gray's 803 that has been the focus of attention for the 803 Working Party.

John Gray was a self-trained printer who developed an abiding interest in computers. He supplemented his knowledge of computers with an Open University degreee in the mid-1970s and justified his continued interest in computers by his concern with computer applications in printing - automatic phototypesetting in particular - a far-sighted pursuit. He told me that he had acquired the 803 to convert 5-hole to 8-hole tape for phototypesetting purposes, though his son, David, recently hinted that printing may actually have been a pretext for his continued involvement with computers.

The personal qualities that I immediately associate with John Gray are his unassuming humility, his exceptional gentleness, the sparkle of his independent and enquiring mind, and his absolute integrity. John Gray was born on 14 September 1919 and died on 28 April 1994.

Book Review - the Leo story

Leo - The First Business Computer, *by Peter J Bird*, 272pp hardback, published by Hasler Publishing at œ15.95 (ISBN 0-9521651-0-4).

J Lyons & Company was the first business in the world to realise that the electronic computer was potentially an invaluable data processing tool, and by quite a long way. As early as 1947, when everyone else involved in computing was interested purely in the high speed computation possibilities, the company decided to put its insight into practice and develop its own computer specifically for business automation.

The result was Leo 1, a first generation machine that was ready for its first practical work in 1951. J Lyons sub- sequently produced two more generations of machines, Leos 2 and 3, both of which were built for internal use but were also offered for sale. The company remained a computer manufacturer till 1963.

Computer historians rarely give J Lyons the credit for either its perception or its achievement in seeing the project through to the point where Leo 1 became the first operational purpose-built business computer. This is partly because Univac, which is often given the credit for this pioneering step, had a very effective publicity machine fully capable of capitalising on the Univac I's descent from Eniac.

Lyons' own attitude also played a significant part. As Peter Bird observes in his new book, "Although a quiet revolution had begun, there was no showroom razzmatazz. Lyons did not see much virtue in publicly claiming a technological first."

Bird is the first person to tell the full story of Lyons' remarkable achievement in book form. He lays the background to his story carefully, showing how J Lyons had become one the UK leaders in office automation and efficiency in the interwar years as a direct consequence of the nature of its business, which involved very large numbers of relatively low value transactions.

This made the company receptive to the possibilities opened up by any new type of automation. "The achievements of Lyons between the wars in the field of clerical innovation undoubtedly helped in the rapid development of Leo and its early acceptance".

It also helped in the pushing through of a pioneering project with no guarantee of ultimate success to its conclusion. This meant overcoming not

only a succession of technical difficulties, which naturally form the large part of Bird's subject matter, but also the supply problems which bedevilled every enterprise in the immediate postwar years. There was even difficult in obtaining paper to write the early reports on!

In addition of course there was the question of the necessary investment, which has so often been lacking in Britain for similarly innovative projects. Bird tells us J Lyons' investment in Leo was £150,000 - £2 million in today's money. The company had to sell an awful lot of teas to recoup that sort of outlay.

The major problem in adapting the embryonic electronic computer to data processing applications was getting the data in and out quickly enough. A large part of Bird's book discusses J Lyons' efforts to develop the necessary peripherals. As well as paper tape and punched cards, Lyons experimented with magnetic tape (largely unsuccessfully) and conducted pioneering work with automatic document readers, leading to the development of the turnaround document.

The reader will also learn how many of the procedures necessary for practical business computing evolved, including the documentation of both operational and programming activity, flowcharting, and the testing and debugging of hardware.

Bird's book is thoroughly researched and very detailed - there are no fewer than 14 appendices, giving details of how the machines worked, the breakdown of components used, the meaning of programming instructions, the numbers of systems sold, and other solid information which will delight researchers and lovers of minutiae. In addition there are thorough biographies of all the major participants in the story.

The book nonetheless remains a highly readable account which will appeal not only to former company personnel but also to anyone interested in the formative years of our industry.

NJE

Editor's Note: A reunion for former Leo employees will be held in London on 30 September 1994. For details see Forthcoming Events.

Letters to the Editor

Dear Mr Enticknap,

As an avid reader of Resurrection, I am writing a brief note to correct one matter of truth and a second of hypothesis from your recent issue.

Firstly, misinformation that is becoming an accepted fact concerns the Colossus paper tape reader vis-a-vis computer tape readers of the early days. For example, Brian Oakley wrote in his Guest Opinion in issue 9: "It is a galling thought that the excellent optics on the Colossus gave the wartime pioneers a better reader than the commercial machines we had to use for at least a decade after Colossus had been put to bed".

The Colossus tape readers, which I have seen, read continuous paper tape loops at high speed without stopping. Due to the software-driven inputting of tapes, the commercial paper tape readers had to start and stop completely randomly as far as an engineer's eye could see. This was a very different requirement.

The Elliott tape reader, for example, stemmed from a design under Maurice Wilkes' aegis at Cambridge, which we at Elliotts made into an engineered production form and supplied him with a few models as a licence fee. This was a very good deal for Elliotts and I am afraid we took rather a long time to provide Cambridge with the engineered models.

These readers were fast and worked using the now well known `rapid clamp to stop' principle. However, they still needed careful adjustment. I well remember one of Strachey's bright NRDC programmers (probably Colin Merton or Brian Millis) writing a horrific test for the Siemens Woolwich 405 tape reader. This went through slowly speeding up start/stop cycles: any resonance was shown up, and it needed a well adjusted reader to pass the test.

I am full of admiration for the Colossus tape readers and there are articles on them in the literature, but I must defend the totally different usage of the start/stop tape reader.

The real sadness to me about the Colossus achievement conerns the logic. This was, of course, built using therm- ionic valves, but they had no real problems of reliability, whereas machines using double triodes were beset with valve deterioration and had to use regular testing and replacement. If only the Colossus circuitry had been made more widely available!

The second - a more emotional but I believe plausible - issue is that regarding Dr Ross handing his hat and coat to John Coales, as stated in Hugh McGregor Ross's lecture and reproduced in your article. I knew Dr Ross well and he was one of the most polite of men. I and Lawrence Clark believe that he pushed his hat and coat towards Coales with words like "Where do I put these?" in his inimitable accent. I feel that Hugh saw but did not hear. I may be wrong, but the facilities in the Elliott canteen were not self-evident in that respect in those days.

I have a great admiration for both Brian and Hugh, but I do not think it is right that such facts or views become accepted history.

Yours sincerely,

Andrew St Johnston
Vaughan Systems
Ware Hertfordshire
26 April 1994

Dear Mr Enticknap,

In his very interesting Guest Editorial in the Spring 1994 issue, Brian Oakley says "I think it was a young electrical engineer named John Fairclough at Ferranti who first developed the good engineering practice required to stabilise mercury delay lines". I am afraid Brian's memory has been letting him down.

Ferranti never made any use of mercury delay lines, nor did any work on them. John did work for Ferranti (at our Bracknell Laboratory) and did make significant advances in delay line technology - but these were nickel delay lines for use in Pegasus and related equipment. I believe that one of his major developments was the enhancement of the magneto-strictive performance by using a bundle of fine strands in place of the single wire that was normal practice. He thus exploited the fact that magneto-striction is primarily a surface effect.

John subsequently transferred to Ferranti Electric Inc (our New York-based subsidiary) and continued with magnetic delay line development which led to some Ferranti lines taking part in early space exploration. He later joined IBM where his technical management skills flourished, and in due course they sent him back to the UK to take charge of all IBM research and

development activity here. Of course it was not long before the British Government had secured his services as a technical adviser.

Yours sincerely,

Harry Johnson
Oxford
27 April 1994

Dear Nicholas,

I wonder if the Computer Conservation Society is aware that this year is the fortieth anniversary of the founding of Britain's first commercial computer company. I am regularly disappointed about how little Leo gets mentioned when the history of commercial computing is being discussed. Unfortunately there is little left of the computers them- selves so there is not much to "conserve".

By coincidence it is also the centenary year of Joe Lyons, the company that was brave enough to go it alone in developing commercial computers. Surely they should get more recognition as pioneers.

Regards,

Peter Byford
Chairman, LEO Computers Reunion Society
Ware, Hertfordshire
1 May 1994

Editor's Note: for details of a reunion organised to celebrate the 40th anniversary of the founding of Leo Computers, readers should refer to the Forthcoming Events page.

Letters Extra - On Identifying Pegasi

In issue 9 of Resurrection we published an article by Chris Burton describing an innovative methodology for determining the identity of the Pegasus currently being restored by the Society's North West Group. He proposed it in an attempt to resolve the question of whether that system is Pegasus number 1 or number 6. Harry Johnson and Derek Milledge have written to us with information that provides insights into this question, and we reproduce their comments here. If any other readers have knowledge relating to the history of either of these systems we would be delighted to hear from them.

Dear Mr Enticknap,

I was greatly intrigued by Chris Burton's article "Determining the Age of a Pegasus" which appeared in issue 9 of Resurrection. To begin with I was at a loss to under- stand how confusion could possibly have arisen between Pegasus number 1, which was assembled not in the factory but at 21 Portland Place and was used in the Ferranti Computing Service there, and Pegasus number 6, which was a standard product sold to Vickers-Armstrongs (Aircraft) Ltd, Weybridge, and delivered to them in 1957. However, some unpublished notes written by Bernard Swann in 1974 provide some clues.

He says that when ICT took over the Ferranti interests in 1963 they sold Pegasus number 1 to Vickers-Armstrongs to supplement their existing machine. In due course Vickers- Armstrongs presented it to their near neighbour, the Brooklands College of Technology, with which they had a close relationship. Bernard says it went on working there until 1969, when Vickers-Armstrongs offered the college the Pegasus II which they had bought in 1961.

It is easy to see how confusion might have arisen between Pegasus number 1 and Pegasus number 6. I have no knowledge about later developments but there surely must be people who could enlighten us.

Chris Burton's ingenious approach should solve the identity question. One possible complication is that manufacture of many of the early boards was carried out by external sub- contractors - tradition has it that one of them was a then little known small company called Racal. It seems likely that such boards would not have followed the same numbering system as those made in-house: but that might make dis- tinguishing between number 1 and number 6 very easy.

A more general observation is a word of warning about using computer lists to identify individual computers. Several such lists were prepared for a variety of purposes and there are discrepancies between them. For example some number Pegasus I and Pegasus II separately and some do not.

Harry Johnson
Oxford
27 May 1994

Dear Mr Enticknap,

In issue 9 Chris Burton considers the identity of the Pegasus at the Manchester Museum of Science and Industry. This Pegasus I has a 4096 word drum and can hardly be Pegasus I number 1, which was modified for the 7168 word drum in 1958 and scrapped about 1965. The Pegasus in Manchester came from Brooklands Technical College, but was said to have been at Vickers Armstrong Aircraft previously. Pegasus I number 6 was the only Pegasus type I at Vickers and it is known to have moved to Brooklands Technical College around 1965: this surely must be the Pegasus now preserved in Manchester.

With reference to Pegasus number 25 it is worth noting that spares for this machine were acquired by University College when others such as number 19 were scrapped; this may account for the secondary cluster of type 01 package serial numbers around 4300.

Derek Milledge
Bracknell, Berkshire
15 June 1994

Internet addresses

Readers of Resurrection who wish to contact committee members via electronic mail may do so using the following Internet addresses.
Chris Burton: chris@envex.demon.co.uk
Martin Campbell-Kelly: mck@dcs.warwick.ac.uk
Adrian Johnstone: adrian@dcs.rhbnc.ac.uk
Robin Shirley: r.shirley@surrey.ac.uk
Doron Swade: d.swade@ic.ac.uk

Working Party Reports

Elliott 401
Chris Burton, Chairman

The Working Party has been in abeyance since the last report in issue 9. The units of the system have been moved to the Science Museum's archive store at Blythe House. It has not been possible to do any work on them there yet, as the room originally planned to house the machine has turned out to be far too small to give sufficient working space. A new plan has been drawn up to occupy a much larger room, which is at the moment filled with punched card machinery. The process of moving these items out, reinforcing the floor and providing power is proving to be long drawn out.

Meanwhile, the 401 cabinets are on pallets packed rather tightly together. As an interim measure, it is intended to spread the pallets further apart. When this has been done, a Working Party meeting could take place to continue elucidation of the machine logic diagrams.

North West Group Pegasus
Charlie Portman, Chairman

The five members of the working party are meeting at the Museum of Science and Industry in Manchester every other Tuesday. Our task is to preserve the machine: it is not intended to restore it to working order at present. We have heard that the project will receive a grant from the Museum and Galleries Commission Prism fund.

We have cleaned the outside of the cabinets and have given the insides a preliminary vacuum clean. We are now con- centrating on cleaning the packages and recording their current positions and state. This work is expected to take all summer. After it is complete, we will move on to the power bay.

Science Museum Pegasus
Chris Burton, acting Chairman

The Pegasus at the Science Museum now stands in lonely splendour in the Old Canteen, with the ghosts of past CCS activities all around it. But the system is in as good health as it has ever been. There have been several "by

request" demonstrations to parties of VIPs: at only one did the machine not work. The members of that party were never- theless happy with what they saw, and the simulator proved useful to show them what might have been!

A highlight was a visit by a television crew to use the background of Pegasus and the simulator to illustrate an item by Doron Swade on preserving historic software for the BBC2 series The Net. Several hours of filming were edited down to about three minutes for the programme.

One Working Party day was held in early June, when five members improved the reliability by eliminating a number of vibration-sensitive valves and connectors. Thereafter the machine ran well for several hours on the engineer's test programs. Some work was also done on the PEGTH tape prep- aration set, and a short training session was held. Thanks to those who attended on a lovely sunny day which they would have preferred to spend at home!

The future of the system remains uncertain - whether we should move it again or not, where it could be moved to, who will occupy the Old Canteen and will the machine be made inaccessible thereby. Efforts continue in trying to find a solution to these problems.

Elliott 803
John Sinclair, Chairman

The main thrust of working party activity over the past couple of months has been the restoration of a second Elliott 803, presented to the Society by Andrew St Johnston to whom we are very grateful. It is not as fully configured as the Science Museum machine, being without film handlers, but is a complete system.

It is our plan to restore this system to full working order. It will then form one of the centrepieces of the Society's museum collection at Bletchley Park. The immediate objective is to have it ready for the formal opening of Bletchley Park, which was scheduled for mid-July as this was being written.

We took delivery of the system from Andrew on 6 May. It had been stored in a barn for a number of years, and was there- fore in somewhat less than pristine condition. On 21-22 May we devoted a full weekend at Bletchley Park to undoing the ravages of the years.

The working party on this occasion included Peter Onion of British Telecom; Dr David Bisset of the University of Kent at Canterbury and his

wife Florin; four of David's students, namely Neil Stevenson, Peter Tilsley, Andrew Webber and Alex Zivanovic; another gentleman whose name I unfortunate- ly did not record; and myself. Each of these volunteers not only devoted a weekend of their time but also shouldered their own expenses.

The work involved taking off all the doors so that we could get at the printed circuit boards. Each one has 60 gold- plated contacts, and each of those needed cleaning on both sides. The work was completed with such vigour and thorough- ness that this 803 is now in cleaner condition than the one in the Science Museum!

Since that weekend further hard work has brought us to the point where we can power the machine up, which is very good progress in such a short space of time. We are now in the process of making it operational.

The Science Museum machine, meanwhile, has been in Blythe House for several months now. Working party activity has not yet resumed, as we await the preparation of a computer room within Blythe House for the machine (this room will also house the Society's Elliott 401 and DEC systems).

The necessary work includes installing a false floor, which is needed for structural purposes (to ensure proper weight distribution) as well as to conceal the cabling. While we wait, the 803 is stored in rather cramped conditions in another room, so restoration work is not currently possible.

North West Group

The North West Group would particularly welcome any material - data, manuals, film - on the Ferranti Sirius and Metrovick 950 computers. Anyone who can help here or who would like to play any part in the group's activities should contact secretary William Gunn at 23 Chatsworth Road, High Lane, Stockport, Cheshire SK6 8DA: tel 0663 764997.

Forthcoming Events

6-7 August 1994, and fortnightly thereafter: Guided tours and exhibition at Bletchley Park, price £2.00

Exhibition of wartime code-breaking equipment and procedures, plus 90 minute tours of the wartime buildings.

30 September 1994: Leo Computers Reunion

This event is being organised to celebrate the 40th anniversary of the founding of Leo Computers as a separate company, and will be held in London. It is open to all who have connections with the company: interested readers should contact Dick Warren on 0753 885063, or at 47 The Uplands, Gerrards Cross, Bucks SL9 7JQ.

4 October 1994: North West Group meeting

Highlights from the best of the London meetings.

13 December 1994: North West Group meeting

Professor Grimsdale will speak on "The transition from valve to solid state machines".

The North West Group meetings will be held in the Science Museum Conference Room, Manchester, at 5.30 pm.