

Computer Conservation Society

Aims and objectives

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

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

The aims of the CCS are to

- ◇ Promote the conservation of historic computers and to identify existing computers which may need to be archived in the future
- ◇ Develop awareness of the importance of historic computers
- ◇ Encourage research on historic computers and their impact on society

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

The CCS is funded and supported by 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 ICL and Vaughan Systems.

Resurrection

The Bulletin of the Computer Conservation Society

ISSN 0958 - 7403

Number 21

Spring 1999

Contents

News Round-Up	2
The Spread Committee <i>Sir John Fairclough</i>	4
The ICT 1301 Series <i>Hamish Carmichael</i>	11
Developing Software for the 1900 <i>Peter Hunt</i>	19
More on the EMI Business Machine <i>Stewart Hine</i>	23
CCS Collection Policy	25
Society Activity	26
Bombe Rebuild Project — Appeal for More Help	30
Forthcoming Events	31
FTP, Web and E-mail Addresses	32

News Round-Up

Dave Holdsworth and Simon Lavington were co-opted onto the Committee of the Society in October.

- 101010101 -

Dave Holdsworth works in the University Computing Service of the University of Leeds. Dr Holdsworth is a specialist in software conservation, covering not only the preservation of documents but also the emulation of historical software on modern hardware platforms.

- 101010101 -

Simon Lavington is planning a new edition of his influential work “Early British Computers”. The Society is helping with this project. Any reader with information which might be suitable for inclusion in this second edition is invited to contact Professor Lavington at the University of Essex (for full contact details see inside back cover).

- 101010101 -

The British Computer Society has honoured Chris Burton with the first of its new Lovelace Medals, which are presented to individuals who have made a major contribution to the advancement of information systems or have added significantly to the understanding of the development of information systems.

- 101010101 -

Tommy Flowers, the engineer credited with the major role in translating the Bletchley Park Colossus concept into reality, died in October aged 92.

- 101010101 -

We have been contacted by a French computer conservation society called Aconit. This organisation possesses a large quantity of hardware, including an ICT 1300 and two Digital Equipment PDP-9s, and has just acquired premises in which to house this equipment. Work has now started on deciding what can be restored to working order and what should be preserved as static displays. Aconit has a Web site at <http://www.digiweb.com/~hansp/ACONIT>: its address is 10bis Rue Ampere, Grenoble, France.

- 101010101 -

Readers who have general queries to put to the Society should address them to the Secretary at the address given on the inside back cover. This may seem obvious, but several people have complained recently that such general queries have gone unanswered; the only explanation is that they must have sent them somewhere else. In particular, addressing queries about the CCS to BCS headquarters is a surefire way of ensuring that they get dealt with belatedly if at all.

- 101010101 -

By the same token, members who move house will not receive any more copies of *Resurrection* unless they notify Hamish Carmichael of their new address. Informing the BCS will get you nowhere, because the CCS membership is entirely separate. It is necessary to maintain two lists because the majority of BCS members are not members of the CCS, and many members of the CCS are not members of the BCS.

- 101010101 -

Just as we went to press we learnt the sad news that Charlie Portman had died at the age of 65. Charlie had been playing a major role in the activities of the North West Group. We plan to publish a full appreciation in our next issue.

- 101010101 -

The Spread Committee

Sir John Fairclough

The Spread Committee was the body set up by IBM to formulate plans for what was to become System/360. This article provides an insider's view of the Committee's deliberations.

The driving force for a compatible family of machines was the programming dimension, which was beginning to look like an enormous investment both by our customers and by IBM. So the desire and commitment to attempt a compatible family of machines was driven not by engineering bigotry but by the desire to have a series of machines that could be optimally programmed to conserve our own and our customers' resources.

The Spread task force was formed to create a development plan. Spread, incidentally, stood officially for Systems Planning Research and Development. Unofficially we made it mean Spalding's Plan to Recognise (or Reorganise — I can't remember which) Each and All Divisions.

In the early sixties, IBM was developing computers in five or six divisions, and anarchy reigned. Individual business units were given considerable freedom to operate in what were described as their self contained markets — but of course they weren't self contained markets. A considerable overlap and duplication was beginning to emerge.

We in Hursley had spent two years prior to the formation of the Spread committee designing a scientific computer and modular processor, called Scamp. We were very proud of this machine. It was the first IBM machine, and probably the first machine anywhere, to embody Wilkes' invention of microprogramming with properly engineered Read Only Memory.

It was at this stage that IBM management realised that anarchy was getting out of hand. So all machines under development were shelved, including Scamp, in favour of the grand plan to be formulated by the Spread committee.

I was invited to join the task force primarily because I was so unhappy with the termination of Scamp that I was threatening to leave. The prize was come and join this task force. I was asked to go and 10 days later my wife and family were on the plane and my house was rented. IBM then and still today stands for I've Been Moved!

The Spread task force sat for 90 days in a hotel that looked very elegant at the time but is now very down market — the New England Motel, just

off the New England throughway. We finalised the report in December 1961, and made our presentation in January 1962.

Members of the task force included as chairman John Hanster, who was leading the division responsible for the 1401, and as vice chairman Bob Evans who was leading the division for the high end machines. John was promoted in the middle of the task force's deliberations, and Bob took over. If you were to identify one individual who would carry the lion's share of the blame, or the accolade for what eventually emerged, it would be Bob Evans.

The task force members had varying backgrounds, and by and large hadn't met before. Intellectually, in terms of the detailed engineering programming architectural arrangements, there was no doubt that Fred Brooks was by a long way the largest contributor.

Systems under development at this time included the 14LC (a low cost 1401); the 7030 (the top end of the line); the 1410 X (a souped up 1410); and the L²C, which was conceived as a small entry machine that would come in below the 1401. In fact it ended up being the model 20.

There was an emerging solid logic technology: it was not yet large scale integration but it was multiple chips on ceramic carriers. That was to be the heart of the planned new machines.

The task force's job was to produce not a design but a vision, plus some underpinning rules that would guide the subsequent design, development effort. One line of processors would serve for all kinds of applications, spanning the range from below the 14LC through and beyond the high end of IBM's existing line.

That's all the management were told about the processor characteristics. Generally IBM senior management in this era had been brought up in the punched card days. They weren't particularly knowledgeable about computer architecture and so our presentation was low on technical content.

There were to be five compatible processors. These were characterised simply by a two-address-add time of 275, 25, 5, 2.5 and 1 microseconds respectively.

Interestingly, our report didn't commit totally to upward and downward compatibility. We all felt it was technically possible, but whether one could produce a low end model at competitive costs and hence a competitive price was the concern. The question was whether the additional cost burden at the low end was going to be affordable in the overall plan. So

we left ourselves a little out. That's the only time this reservation ever got expressed: all future presentations assumed 100% upwards and downwards compatibility.

One recommendation was that decimal digits should be represented in four- and eight-bit bytes. The biggest argument that we had during the whole of the implementation was whether we wouldn't be better off with a six-bit byte, simply because that was cheaper. I would say that there was more blood spilt on this argument than on any other. It had to do with the cost burden of the low end of this series, carrying the full compatibility characteristics. That became so much more easily achievable if we had a six-bit byte.

We recommended that each four-bit byte was to be directly addressable. IBM never implemented that, not even at the high end.

Other recommendations were that negative fields should be represented in true, not complementary, form; and that address modifications through added indexing should be standard across the series. The Spread committee also specified that memory protection should be standard.

Floating point was to be available at least as an option. One of my contributions was to specify sterling arithmetic, in pounds, shillings and pence! There were to be uniform subroutine linkages, and program interruption from external signals on all machines. There was to be a real time clock.

I/O was to be programmed through channel control words, whether there was a discrete channel implemented or whether the channel was integrated into the base logic of the machine. I/O operations were to be overlapped. There was to be cross channel switching. We recommended an integrated front end for terminal attachment, which in the event was never given serious attention as an integrated part of the design — that came later. The range should support any I/O device type, and to any I/O device type all channels should appear identical. When one processor was substituted for another the I/O should not be changed.

All data parity was to be checked. The maintenance characteristics of the machine should also be compatible. And we specified multiple cpus for ultra reliability.

My major contribution to the debate was on microprogramming. Microprocessor controls were to be used throughout the series except, particularly for the high end machines, if the designers could demonstrate that they could produce a classical hardware control system that was less than

two thirds the cost of the read only memory microprogramming solution.

There was to be a single architecture for memory-to-cpu coupling, a standard interface. Many of these recommendations seem so self evident today but in 1960 they were not, and these ground rules turned out to be very pivotal in what subsequently developed.

That is not a comprehensive list of the engineering ground rules, but it highlights the key points that were in our report.

In our presentation to management, we put forward our vision that we could minimise the impact on existing installations and existing customers by emphasising new applications, while through compatibility allowing customers to grow with minimum disruption. We would have a unified source language, or languages, to give us a more efficient way of producing our own application software as well as system software.

On the marketing front, we urged better pretesting, so that all customers would be influenced by this announcement. There was a great deal of concern about keeping our plans confidential, especially the fact that we had conceived a compatible family. We went to great lengths to protect that concept, even though the plans were initially to dribble the machines out one by one rather than in one big bang. We envisaged that one day, after two or three announcements, we would declare that we had a strategy for a compatible family.

At the time we were deliberating these issues, we were in many ways putting our fingers in the breeze: our confidence wasn't 100% that we could ask for a massive investment by management on the basis that we'll announce all of these machines at once and the following will result. We weren't sufficiently confident to say that: we needed more time to get the design detailed and finalised.

Our presentation satisfied IBM management's requirement for a 20% per annum growth rate. When Fred Brooks retired, he left IBM to become a professor of computer science at University of North Carolina. He subsequently made a speech to IBM management in which he pointed out that if IBM continued growing at the current rate of 15% or 16%, at some point the company's turnover would exceed the gross national product of the United States. Keep on growing and it would pass the gross national product of the world. It seems to me, he said, looking Tom Watson straight in the eye, that the name of the game in this company is not to let it stop on you. And it did! But not on Watson!

We then put forward a number of comments on how all this could be made to happen. I've described how IBM had at the time a number of business units all satisfying their own computing requirements: there was a need to pull this all together managerially. What was proposed was some centralised control of architecture, so that while processor development would be undertaken in existing locations Bob Evans, who ended up as project manager for the whole series, could be sure that he had the authority to control these disparate groups.

So a corporate control organisation was created. Gene Amdahl became the Chief Architect, reporting to Fred Brooks who in turn reported to Bob Evans. The development locations such as Hursley and Endicott wore two hats. I continued to report to IBM World Trade but was also part of this organisation, so there were times when life became quite rather stressful. Budgets and schedules and plans were provided by—in my case—the IBM World Trade organisation, to satisfy demands of this centralised plan. It became tense from time to time but we managed it.

Brooks' approach was to assign pieces of the architecture to the implementation groups to define, on behalf of the whole family. That was a very effective technique. It had the effect of building a team out of these disparate groups. I remember we in Hursley defined the fixed and floating point arithmetic set for the whole family. I think Ivor Jones did it. So the team effort was strongly emphasised by Fred Brooks and it worked very well.

Now we come to some of the technical characteristics. Main memory had cycle times going down to half a microsecond. It seems like forever today, but that was the fastest core memory that was envisaged with the capacities we wanted. Read Only Memory was to have cycle times a quarter to a tenth of the main memory.

The most demanding targets were on cost: generally, components had to cost a tenth to a half whatever they were costing at that time.

IBM had a richness of technological capability and all of it could be marshalled to support this effort. So it was a brand new design, a brand new everything really. Every piece of technology was devised specifically to serve this family. So an important part of our report was to marshall those technology resources in IBM and give them some objectives.

We wanted binary addressed files. The Rmac was not binary addressable, and a number of subsequent disc files were not binary addressable as a result. The tape group's targets were 500 characters a second read/write

speed and \$200 monthly rental. IBM was mainly renting equipment at this time, so the rentable value of hardware was the way that cost/performance was primarily measured then.

The closest we got to considering a multiplexer was a cross point switch or a time multiplexer for switching between I/O channels. That did not get done in a timely way. It was subsequently done very well but as a second stage.

Then to give the myriads of staff members in the corporate office and the divisional offices something to do, we devised schemes to measure the progress of this overall effort.

We began to take 100% up and down compatibility for granted more as time went on, and any doubts that we might have had were subsumed in the thought that this quality was so vital and would be such a valuable contribution.

There was no discussion of procedural languages in our report other than as a throw away line: if we can realise a compatible series of hardware, wouldn't it be nice to have a NPL (New Programming Language) that would play a complementary role. This didn't get the serious attention it deserved until two years later.

In marketing terms, the project was seen as an opportunity for new business and new applications. "Impact" was a word on everybody's lips at IBM in those days: introducing a new product would "impact" (ie displace) existing installed products, so how you managed and controlled the smooth flow of revenue and then realised a growth of 20% per annum was a very important business judgement.

The marketing people also made some technical suggestions: for example, they thought we could use the 1400 as an I/O controller effectively for this new family. Architecturally that was never on.

Another thing we envisaged was a complementary concerted effort across the disparate business units, to look at industry application development. Some straightforward direction was applied in order to get that marshalled. I would say it wasn't wholly successful. There were still many different views about the unique requirements of individual vertical and geographic markets, and to get everybody singing out of the same hymn book proved to be more difficult than anybody imagined.

Introduction planning in the report envisaged a staggered announcement. Machines number two and five were to be announced first. In fact my team built the first 360 machine in Hursley. After development started

there was pressure for what became known as “tiny biny” — a subset of the model 40 aimed at the medium scientific market. We got committed through that mechanism to an earlier schedule than the other four processor development groups, and that was why we were first.

The grand plan was announced on 7 April 1964. At that point the bulk of the engineering and of the architecture and system testing had been done on the machine that we built. The other machines in Poughkeepsie and Endicott were some months later and weren’t as advanced as ours.

Initially the Spread committee planned no announcement until processor one reached a test lab. Processor one was the small one and it was this compatibility point — can you actually get competitive cost performance with a full and comprehensive set of functions at the low end of the line? — that led to the desire to have it thoroughly tested, with enough programming to make it run.

L²C was the entry-level machine IBM management envisaged before this whole project started: as I said, it ended up as the model 20. That embodied just a subset of the full System 360 functionality, although the Spread committee had argued that it should adhere to the full architectural spec.

We argued that the low cost 1401 should go ahead because that was close to being ready anyway, and would buy us some time until we could get the family announced. Similarly, we recommended that the souped up 1410 should continue. At the high end of the line, we thought the 7030 should not be announced until there was enough known about the high end 360 machine — processor five — to decide whether we could afford to go with just the 360 machine.

Finally, we also considered security. We introduced the “registered IBM confidential” procedure, and every project document was numbered and tied to individual names. It became very tedious, but by and large it worked during that period.

Editor’s note: this is an edited version of the talk given by the author to the Society at the Science Museum as part of the IBM System 360 seminar on 21 November 1995.

The ICT 1301 Series

Hamish Carmichael

The first computer developed by International Computers and Tabulators (ICT) after its formation in 1959 was the 1301. It became the mainstay of the company's computer product line until the introduction of the 1900 series in 1964. In this article the Secretary of the Society describes the machine, on which he worked himself in the early part of his career.

The merger which formed ICT was formally announced on 19 January 1959, and as with all the subsequent mergers until well after the formation of ICL, the company immediately found itself with incompatible ranges of computers to make and sell. The Hollerith (BTM) side had delivered respectable numbers of the 1200/1201/1202 series, but the design was already seeming old-fashioned. Powers Samas had not had comparable success with the PCC (which many BTM people said wasn't really a computer at all). There was also at the time a potential marketing arrangement with Ferranti to sell a Ferranti machine — perhaps a Pegasus variant? — under the name Pluto.

Faced with this situation, ICT decided to set up an independent computer design company as a collaboration between ICT and GEC. From GEC's point of view, a stimulus was that the company had just lost a big and important order through not having any computers in its repertoire, and it was very keen to make up that shortfall.

The result was called Computer Developments Ltd, normally shortened to CDL, and it set up a small design office initially in a nissen hut at GEC Wembley. Later it was moved to Kenton, and still later to Harrow. It is said that GEC thought ICT knew all about the design of computers, and that ICT thought GEC knew all about the manufacture of electronics, and that both parties were wrong!

Most CDL people were ex-GEC, apart from Raymond 'Dickie' Bird and Harry Foster from ICT. John Wensley and Dickie Bird were the chief ideas men, and produced project outlines P1 to P5, one of which, P3, became the 1301. CDL was to be the system and logic designer, but with no involvement in physical development or manufacturing. The prototype, with a one microsecond pulse rate, was to be developed at GEC Research Labs at Wembley.

In the event, ECH Organ (inevitably and invariably known as ‘Echo’), later the ICT Director of Manufacturing, quickly came to distrust the Wembley build, and developed a parallel prototype at Coventry under Colin Ladds and George Gibson. This became the basis for the production machines.

This design was notable for the extensive use of wrapped joints, chosen for their reputed reliability. Some still swear by them, but more seem to swear at them — there are greybeard retired engineers for whom the 1301’s wrapped joints have remained the stuff of nightmares. The choice of this technology apparently also had a major effect on the physical layout of the major component units — without the use of cables, the units had to abut each other in a precise sequence.

Other people involved *ab initio* included Dennis Espley, Norman Bligh and John Beazley. They were joined up to a year later by David Lush, Bob Whittall, Fred Dearnley, Sean O’Leary, David Wood and Jack Cotton. Steve Shirley and Kevin O’Brien were involved in early software development.

One very successful innovation was the involvement of Noel London, an experienced expert in industrial design, who from day one had an authoritative influence on the ergonomics of the design, including the colour scheme, the standard height for all the units and all the details of the final appearance of the covers, the indicators, and the manual controls. The 1301 was always an extremely good-looking system.

Dickie Bird recalls that the CDL design offices were fortunate to have a large (15’ by 50’) conference room. For testing purposes, the logic diagram of the whole machine was drawn out on the floor, and volunteers representing pulses walked from gate to gate. When two people tried to stand on the same point at the same time, or one tried to go through a gate that wasn’t open, it was a good indication of a logic problem. The whole thing was very humorous but very effective.

The front row of the standard configuration of units comprised, reading from left to right, the card punch, card reader, console and line printer. Behind the console stretched the spine of the machine, containing the processor. At the rear another transverse range contained drum(s) to the left and core store units to the right. Magnetic tape units were not part of the integrated configuration, so could be placed anywhere convenient. The same went for paper tape equipment.

Two card readers were available, operating at 300 or 600 cards per minute. A single card punch had a maximum speed of 100 cards per minute. Three line printers offered 80 print positions and 300 lines per minute; 120 print positions and 300 lines per minute; or 120 print positions and 600 lines per minute. A paper tape reader could be attached, offering a maximum speed of 1000 characters per second. A paper tape punch could operate at up to 300 characters per second.

In the relatively short life of the series no fewer than three ranges of magnetic tape units were approved for connection to 1301 machines. Their specifications covered: one inch tape, with transfer rate of 90,000 digits per second; half inch tape, with transfer rate of 22,500 digits per second; and quarter inch tape, with transfer rate of 16,500 digits per second. Looking back, it is surprising that such variability was thought viable; it would seem that the numbers would always have been too small, and the margins too squeezed, to justify it.

As a slightly more exotic attachment there could be an interrogating typewriter, though I believe few were ever implemented. An input message would be typed in red as it was being keyed, and the program's response would appear on the following line in black. It is surprising how quickly the red/black capability of bi-coloured ribbons, which had been successfully exploited for a long time in the tabulator era, disappeared from the computer world.

The basic unit of storage was a word of 48 bits arranged as 12 digits. With one exception, binary was not used internally, all operations being carried out on numbers held in character form. The fact that a character could be composed of bits valued 1, 2, 4, and 8 helped with handling pence, from 0d to 11d, as a single digit. The choice of binary-coded-decimal was dominated by the observation that input/output performance was more important than sheer calculating speed in all commercial working.

The Immediate Access Store (IAS) consisted of modules of blocks of ferrite cores representing 400 words. A system could have from one to five modules, giving a capacity range of from 400 to 2000 words. (Later on, the 1302 could have up to 4000 words of IAS). Each word was directly addressable, so the basic address range ran from 0 to 1999.

The second level of storage was on a magnetic drum, with a capacity of 12,000 words. Up to eight such drums could be attached to a single system. For reduced configuration systems the drum could be supplied with a cut down capacity of 6000 or even 3000 words; in such cases the

actual drum was the standard component, but the number of read/write heads was reduced appropriately.

Drum storage was arranged along its length into 60 channels, each channel having 200 words. A channel was further divided logically into ‘decades’ of ten words each. A single decade was the minimum unit of transfer in either direction between drum and IAS.

It was very common—and easy—to transfer a complete channel at a time to or from an IAS location which was a multiple of 200. In channel transfer mode, the transfer would begin with the next decade to come under the read heads, which was very efficient. Alternatively, any number of consecutive decades could be transferred, and in this mode the transfer could overlap from one channel to another. Clearly in decade transfer mode the transfer could not start until the first specified decade came round to be read.

The arithmetic unit consisted basically of three 12-digit registers—A, B and C—and a ‘mill’ where actual additions and subtractions were performed. The function of the registers was as follows.

- A had a 48-bit parallel connection with IAS, and was therefore used for: all store accesses; reading and writing data between IAS and registers B and C; and reading instructions from IAS to be obeyed. It was also used for all transfers to or from the drum and to or from magnetic tape.
- B was involved in most arithmetic operations and was also the link between the central processor and line printer, card punch, paper tape reader, paper tape punch and interrogating typewriter. It was also required for all operations involving left or right shifts.
- C was used during multiplication and also for communication with the card reader.

The use of the registers during peripheral transfers needs some further explanation. It may seem very strange to the modern mind, but the peripherals on the 1301 series were unbuffered. So instead of just saying ‘Read a card’, the programmer had to initiate the movement of a card, and then by program test for the arrival of each six columns in Register C, and then store them before the contents of the next six columns could arrive (six columns rather than 12, because each column was represented by two components—the zone and numeric parts).

Similarly, when a line of information was to be printed, the outward transfer was controlled by program in precise detail. As each row of characters on the print barrel approached the printing point, the program had to scan the line of information in store, and in a process known as ‘row binarisation’ set up 120 binary digits with ‘1’ at all positions where that character occurred and ‘0’ in all other positions, releasing this to the printer at the exact moment when the print hammers were due to fire. This was repeated in turn for each of the characters in the printer’s repertoire.

This approach to peripheral control was a commercial decision. Because at the time storage was so expensive, in order to keep the total system cost down it was judged better to get the computer to do the work directly.

Standard modules of program code for each type of peripheral transfer, handling all the niceties of timing and incorporating all the necessary testing for error conditions, were of course supplied in the manuals and on training courses, so that few application programmers had to grapple with such low-level problems unaided.

And, since many early applications mimicked the processing cycles of the preceding generation of punched card equipment, it was soon decided that the programs for peripheral transfers could be interleaved and synchronised, leading to the development of ‘PPF’ (Print, Punch and Feed), one call to which could print a line, punch a card and feed a new card, or any subset of these operations. It was said at the time that there was a suitable analogy in Hoover’s then current motto: “It beats as it sweeps as it cleans”. Later a higher level of abstraction was provided by the PPF Control Routine. For the ingenuity of these procedures much credit is due to Geoff Cuttle.

Most instructions were six digits in length, comprising a two-digit function code and a four-digit operand. Two instructions could therefore be fitted in a single word. There were a few double-length instructions, both halves of which had to be held in the same word. Only the first instruction in a word could be the destination of a jump instruction. These restrictions meant that the second half of a word was sometimes left empty—effectively being occupied by a ‘Null’ instruction.

Numerous indicators could be tested by program. Automatic indicators covered conditions during arithmetic, showing positive, negative or zero, overflow, and all the details of peripheral transfers, including transfer errors and timing errors. An automatic indicator that was permanently

'on' provided an unconditional jump instruction. Programmed indicators could be set, unset and tested by program, entirely under the programmer's control. Thirdly, 10 manual indicators, settable by switches on the console, provided another way of varying the operation of a program.

The engineer's test program for the manual indicators ensured that they were always in tip-top condition. A glamorous lady was printed on the line printer. As each switch was set 'on' she shed part of her clothes, until when they were all 'on' she was all 'off'.

Apart from the necessary switches to control the operation of the machine, the console had rows of neons to display the contents of registers A, B, and C, the contents of the control registers CR1, CR2 and CR3 used in the execution of instructions, the state of programmed indicators, and the state of transfer error indicators. It was possible to set numbers in the registers and control registers manually, and step through a program, one instruction at a time, under manual control.

A simple form of relative addressing was used to give programming sufficient freedom from absolute addressing to ensure the necessary flexibility for easy writing and easy testing. Each block, whether of program code or data area, was given a distinguishing number, called a Relativiser Reference Number (RRN).

Thus when coding the programmer could say 'Transfer Register B to the 4th word of data area 17' without needing to know precisely where that word would occur at run-time. Similarly, 'If the result in the Mill is negative, jump to the first instruction in word 25 of block 8'. RRNs were converted to absolute addresses, and the subsequent relative addresses were correspondingly converted, when a program was read in under 'Initial Orders'.

I vividly remember, as my first venture into serious programming, writing the central time analysis module of 1301 Pert in the basic assembly language in 1962/63, shortly after ICT had occupied Bridge House North, its first foothold in the Putney (strictly Fulham) area.

The input module and the far more complex output module were written with great skill and economy by a charming and brilliant young fellow Scot, John O'Rourke, who was tragically killed in a road accident shortly afterwards. Our work was subsequently converted into 1900 Pert, in which form it had a long and very successful life.

There was also a somewhat more advanced programming language known as TAS (for Thirteen hundred Assembly System), which existed

in several versions. Thereafter there appeared an early version of Cobol, and associated with it a scheme called 'Rapidwrite'.

This provided 13 pre-printed and brilliantly colour-coded card types, each card type covering a simple function such as Read, Write, Move, Perform, Compute. The pre-printed surface of the card could be used by the programmer to write the necessary details into boxes of the right size and in the right sequence. The punch operator read these details and punched them in the appropriate columns, which were indicated by further printing along the bottom of the card.

Even before the 1963 acquisition by ICT of the Computing Division of Ferranti, the Manchester Autocode used on, I believe, Mercury and Pegasus was also implemented on the 1301 series. Using this, my second major programming effort was an attempt to automate the setting of the 'Serial', the schedule which reconciled ICT's manufacturing capacity with sales forecasts and supposedly produced a mutually agreed production programme.

(The project was doomed from the start; it failed to take account of the amount of human understanding, demonstrated by such practitioners as Henry Harnack and Mike Forrest, needed to adjust for the optimism—or pessimism, but usually optimism—of sales forecasters, to incorporate knowledge of unannounced future products, to allow for competitive activity, and so on.)

A *Computer Survey* of June 1962 lists 34 orders for 1301 systems, starting with British Railways Eastern Region, Rubery Owen & Co, South West Regional Hospital Board of Bristol, Cadbury-Fry-Pascall of Hobart, Tasmania, Scaffolding of Great Britain, British Petroleum, Esso Petroleum (Ireland), London University, South African Railways and Harbours—three systems in Cape Town, Johannesburg and Durban, and so on. At number 20 in the list is the Liverpool Victoria Friendly Society.

The machine that went to Hobart was later transferred to Dunedin in New Zealand and remained there in service at Cadbury's factory until it was superseded. It has fortunately been preserved, and more recently restored to immaculate condition, with a hope of restoring it further to working condition, by Bruce McMillan, a member of this Society. Any member visiting New Zealand is strongly encouraged to go and see this magnificent machine, along with an interesting representative collection of Powers Samas 40-column equipment, in the Otago Settlers' Museum, Dunedin.

In October 1998 a member of the Society now living in south-east France reported that, although he had not yet seen it, in that region there is a 1300 system in store, together with a large number of relevant manuals—in French, of course.

Two further 1301s were run until only a few years ago by a group of enthusiasts in Surbiton, operating as Galdor Limited. They are now owned by a member of this Society in Kent, who has, in addition to the machines, an extensive collection of manuals, drawings and spares. All units of these machines are known to have been working within the past few years, though not necessarily together.

One is serial number 6, known as Flossie, built to prototype spec and reputedly the first one to leave the factory, when it was installed at London University. The other is serial number 75, known as Arthur, operated as a spare machine alongside Samantha by the Liverpool Victoria Friendly Society. It is hoped that during 1999 one of these machines will be transferred to Bletchley Park, there to form the centrepiece of a display representing a typical 1960s computer room.

Several members of the Society with 1301 experience recorded in their membership details have expressed a continuing interest in these machines. We hope that enough 1301 expertise survives to enable the machines to be restored to operable condition.

In retrospect, the 1301 series can perhaps be seen as a brave attempt rather than an outright success. ICT could not manufacture them fast enough, which perhaps led to the introduction of the RCA 301 (ICT 1500) as a stopgap.

Also the basic architecture, despite its adaptation down to the 1300 and up to the 1302, did not seem flexible enough to cover the potential range of requirements. And the unbuffered peripherals, with their harking back to the card per cycle characteristics of an earlier generation of systems, could not possibly compete with the Standard Interface of the Ferranti Packard 6000. Therefore the 1900 Series was the logical next step.

Nevertheless, the 1301 sold in respectable numbers—the total may have approached 150 systems. And personally, I continue to respect the 1301 and remember it with considerable affection.

Developing Software for the 1900

Peter Hunt

In the early sixties techniques and strategies for the development and support of software had not yet evolved. So when in 1964 both IBM and ICT announced the first compatible computer ranges, the two companies had to work out from scratch what software they were going to produce and how they were going to develop it. In this article the man responsible for ICT's software production tells how he tackled the problem.

In late autumn 1964 I was asked by Arthur Humphreys to take charge of the production of the software for the 1900 computer range. On reflection, I am not sure the word "software" was used: nowadays that's what it would be called.

At the time I was Head of the Bracknell Laboratories running several interesting hardware and software projects, with large penalty clauses. I had about 12 years programming experience, and had learnt the hard way that software projects were easy to conceive but difficult to implement, were invariably late on delivery and usually exceeded budget.

I asked what software had been promised for 1900, and was given a copy of a document — I think it was document number 1024 — which listed what would be available and by when on the hardware of the 1900 as it was envisaged at that time. I have never found out who wrote the document. It was a phenomenal list of software on various configurations including all types of different peripheral equipment.

After reflection, I decided that I would only accept the appointment under certain conditions:

- I would be allowed to revise document 1024 with more appropriate delivery dates;
- I would be given the appropriate resources to do the job, including adequate competent experienced programmers, and adequate computer hardware facilities under my own control on which to develop the software; and
- that I report directly to Arthur Humphreys (so that I could be sure that if I had problems they would hopefully be sorted out quickly without going through layers of management!).

Arthur accepted these conditions, and then the real work began.

First, I rewrote document 1024 and reissued it with more realistic delivery dates. This was not difficult in itself, but it did involve re-educating the sales force, which had been working with the previous edition. I had to spend time in meetings with senior members of the sales force, and give lectures to the junior members, assuring them that the new dates would be kept, and that there would be no further slippages during the lifetime of the project.

Second, we started work immediately on recruiting 100 programmers.

Third, it was arranged that the necessary hardware would be delivered to “Programming Division” (as we were then known). This included every model of processor and at least one of every type of peripheral that was likely to be delivered to customers.

I organised the production effort into a number of divisions responsible for specific software areas:

- operating systems
- compilers (scientific and commercial)
- general purpose software (eg Plan assembler language, housekeeping software)
- applications programs
- services (including supply of computer time, quality assurance and issue of software).

So there were five divisional managers. I met them every Monday morning to sort out problems and to consider progress on all matters.

At the time, software engineering was either non-existent or in its infancy. The divisional managers decided how each of their projects should be organised, how they would be implemented (there was no BS standard in use) and how they would monitor progress against the promised dates (Pert was in fact one of the applications programs we were implementing).

One of the aspects of the project that worried me most was the problems that might arise when we started to issue all this software to customers all over the world. The originating programmers would have tested it to the best of their ability but, as we all knew, the software would still contain bugs when issued. We wanted these reduced to a minimum.

We therefore set up a primitive quality assurance group. Its task was to take software from the production divisions once they said it was ready for release and, using only the documentation available to customers (we had a separate group, not reporting to me, of technical authors producing manuals and user guides), to use the package as fully as possible, imitating as far as they could the day-to-day usage by customers.

When they discovered errors, the package was referred back to the originating division for correction. The quality assurance team then re-tested the corrected package, and this iterative process continued until the team was satisfied with its correctness. Only then was the software released.

It was not long before 1900s were being delivered in quantity all over the world, and the software distribution system had then to be organised properly.

We decided to issue the software on magnetic tape as the general rule. This meant we had to have fallback arrangements for those installations which did not have magnetic tape transports, such as the 1900s which used cassette tape.

Once issued each software package had to be supported, as we would now put it. When a customer reported an alleged error, this call (or more likely a letter) would be dealt with initially by a front line support force in another division (not reporting to me). This was because in many cases the customer simply needed assistance in the use of the package.

When the support team thought the customer had discovered a genuine error they referred it to us. We investigated and reported back. If it turned out to be an error on our part a “software notice” would be issued to all customers with the following information:

1. description of error;
2. when it would be corrected (release number and date due);
3. what to do in the meantime.

This all seems very standard today, with our modern help lines, but 30 years ago we were breaking completely new ground.

As time went by, more models of the 1900 were produced, requiring more software, while the sales force was also asking for new software products to deal with the competition. By the autumn of 1968, there were over one million words of 1900 software code; by then we had over 1000 customers, located all over the world.

Space to house all the people in the division (now renamed the System Development Organisation) and all the computers under our control was at a premium, at Putney in particular. So it was decided we would move the whole organisation to a new special purpose building to be constructed at Bracknell. This was to consist of an aircraft hangar type building to house all the hardware, and an adjacent office block for all the staff. Initially known as “Hunt’s Folly”, it became one of the largest software production facilities in Europe.

I was particularly pleased when at this time, 1968, ICT was awarded the Queen’s Award to Industry for Technical Innovation. (Many people asked me what was the technical innovation — my reply was simply that we *actually* produced some *software*, and supported it, and that most of what we promised was on time.) I must add that the Award covered not only the software work for which I was responsible, but also all of the Executive work carried out at Stevenage and West Gorton.

It has never failed to amaze me that here we were in a highly technically difficult interface with three groups of programmers working in three separate locations — West Gorton on large systems, Stevenage on small systems and ourselves in Putney producing software for both — and I did not once have to sort out any difficulties with the interface with the Executive teams. Full marks to the Compatibility Committee under the chairmanship of Bruce Paterson.

Editor’s note: this is an edited version of the talk given by the author to the Society during the ICT/ICL 1900 seminar at the Science Museum in May 1996, prepared from a text supplied by the author. Peter Hunt sadly is no longer with us: I am grateful to his former colleague Virgilio Pasquali for checking this edited version for accuracy.

EDSAC99 — 15-16 April 1999

On these dates the Cambridge Computer Laboratory will be celebrating the 50th anniversary of the Edsac 1 computer. For details see the Laboratory’s Web page at <http://www.cl.cam.ac.uk>, or e-mail to edsac99@cl.cam.ac.uk, or phone 01223 334 600 and ask for edsac99.

More on the EMI Business Machine

Stewart Hine

A copy of *Resurrection* number 16, featuring Ron Clayden's article on the EMI Business Machine, has been passed to me by a colleague. I was a development engineer on that project and its pilot, having been transferred with some others from Test Gear Division when that was wound up.

The drum was the heart of the machine, and a fine piece of precision manufacture it was. We engineers were strictly forbidden to touch the heads. Setting them was a delicate business, performed by Alan Stone who visited us about once a week from Wells. A setting tool was clamped to the head support ring and had jaws which gripped the protruding 'tail' of the head. Two screws allowed adjustment of the head radially and in azimuth until a satisfactory signal was obtained. The static clamp would then be re-tightened and the tool carefully detached.

The 16x16 head shifter was something of a trial. Due to unforeseen flexibility in the structure, certain head moves would result in the heads kicking forward and striking the drum, which would ring like a bell, giving a new meaning to the term 'clanger'! Alan would be summoned from Wells, and after repairing the scar with a little brush and iron oxide paint, would reset the head shifter.

Mr Clayden omitted to mention another read-only track — the Initial Instruction Track, on which was written what we would now call the Bios. There was also a Word Pulse output, but whether that was derived from its own track or by division from the clock, I cannot now remember.

The valves were almost all ECC81 twin triodes, with just a few high slope pentodes (6CH6) as buffers. Germanium and silicon diodes were not yet available, and decoder networks were realised using miniature selenium rectifiers. They were fast enough for the modest clock rate, and their large forward drop was not a problem in view of the high signal levels.

Since virtually the entire circuitry was based on long-tailed pairs, at any time about 50% of the triodes in the machine were biased to cut-off, some for quite long periods, and this caused the 'cathode poisoning' and the loss of emission referred to by Mr Clayden. Deteriorating valves were detected by a marginal test in which the bias on the non-signal grid of each pair could be varied by several volts above and below ground, first globally, then chassis by chassis and finally stage by stage. This process was carried out every morning before starting work.

We also had a ‘calisthenic’ program designed to exercise every long-tailed pair in the machine, and later we were able to obtain some special quality valves with improved cathode coatings. These measures reduced the problem to a manageable level.

The transistor specialists had their problems too. The core-logic elements were made up as little potted blocks about the size of a postage stamp. I well remember the consternation when a new batch of these modules, sent up from Hayes, refused to work. It transpired that the original development, using the then available transistors, relied to some extent on hole storage effects. When improved manufacture reduced hole storage, it radically altered the characteristics of the modules!

My main contribution to the machine was the introduction of a totally non-standard valve — a 6F33 pentode designed for time-base circuits. When we started testing the Multiply and Divide functions (not included in the pilot code and hardware), we soon realised that the detection of arithmetical errors demanded a delayed sync facility so that we could examine the process word by word. I designed a suitable circuit, recalling the Sanatron precision timer from my radar days, and was grilled for an hour or so by Ron and Bill Ferrier to prove that a) it was necessary b) it couldn’t easily be done with twin triodes and c) I fully understood how the circuit worked. I managed to carry my point and the delay unit was duly incorporated into the machine.

When the time came to deliver the machine, I put forward the idea of transporting it by canal. This seemed to me to be entirely practical, since it would fit nicely into a narrow-boat, it would be well protected against shocks and vibration, and both EMI and BMC had wharves on the Grand Union Canal. My suggestion wasn’t taken seriously, and the machine was eventually delivered by road on a low loader, fortunately without damage.

The technical literature was not the only thing that was not ready on the delivery date. The machine itself was far from complete, and a team of engineers spent about a year on site completing and commissioning it.

Test Gear Division had offered a number of small projects which a junior engineer could handle with minimal supervision, and I found the total absence of similar work in the computer project irksome. So once the machine was up and running I was glad to escape to the embryonic process control group under Jim Hosking, and later to the Automation Division into which it was absorbed.

Here I worked on CNC machine tools, programmable stacker cranes,

and the EMI Robotug, with which I was associated more or less from start to finish. However, the practice of delivering on a salesman's promised date and completing on site followed me. It became almost standard practice in Automation until the engineers staged a 'peasants' revolt', and refused to go on site for an open-ended period until the equipment had been properly tested at Hayes.

It is customary to end reminiscences like this with the statement that I wouldn't have missed the experience. I would not go so far in this case, but I am grateful for the introduction to a degree of rigour in design which has served me well in my later career.

CCS Collection Policy

The Committee of the Society has formulated a policy statement concerning procedures for dealing with computers of historical interest that come to the Society's attention. This is published in full below.

1. The Society has no Collection of its own, and no premises in which to house one. There is no intention to change this.
2. When the Society hears of historic equipment which is becoming available for conservation, it will attempt to find a suitable home for it in one of the following major collections:
 - The Bletchley Park Museum Trust
 - The Science Museum, South Kensington
 - The Museum of Science and Industry, Manchester
3. The Society will also alert other collections to the availability of surplus equipment, where the major collections are unable to offer to house it, if it fits the appropriate area of interest. Members who know of such collections are asked to ensure that the Secretary is aware of their location and subject matter.

Society Activity

Bombe Rebuild Project

John Harper

Our appeal in the Spring issue of *Resurrection* (number 19) was very successful in recruiting AutoCAD draftsmen. We now have a team of four redrawing the original drawings required for our rebuild. To start with progress was a little slow, but as experience was gained the output rate increased dramatically. The drawings for the central mechanical core of the Bombe are now complete.

Not only are they complete in terms of individual drawings, but also in full assembly drawings. What we have done is to construct the machine on paper, checking every part for fit as it is assembled with the other. This has confirmed that the correct part is being used and that it is accurate.

A few 1990s errors have been corrected. Interestingly, mistakes made back in the early 1940s have also been discovered. Fortunately these were easily understood and rectified. The result is that we have a more accurate set of drawings than was originally available. However we do have far better tools available to us these days, so this is to be expected.

We estimate we have completed over 70% (around 600) of the drawings required in total, and over 90% if one counts only the mechanical parts. The electrical parts such as looms are to be tackled after the mechanical parts are complete. At the present rate of progress I estimate that the redrawing exercise will be close to completion in the spring. However with the core completed we have high confidence in this area, and production can proceed as quickly as the necessary resources become available.

One of the major advantages of redrawing using CAD techniques is that, in some areas, we are able to take advantage of Computer Aided Manufacturing techniques. We have had some plates cut on a laser cutter, and others punched to shape using a large computer-controlled punching machine at Nortel Harlow. From drawings produced by the team we have had a complex eight part casting pattern made at the Warwick University/Rover Rapid Prototyping & Tooling Centre. The casting, incidentally, is required to house the main Bombe gearbox.

We have been very successful with our appeal for Hollerith/BTM parts. We have an example of every part required but are still short on quantity. To those involved, please keep up the valuable work.

Elsewhere in this copy of *Resurrection* is a further appeal for help. We need all the help we can get now, so please come and join our enthusiastic group of helpers.

Elliott 401 Working Party

Chris Burton

The plinth sections have been aligned and bolted together, which now allows the six cabinets of the machine to be similarly aligned and bolted, in readiness for soldering all the inter-cabinet wiring. We have yet to document a procedure for this work, which needs to be well recorded because of the significant physical intervention.

Work has started on replacing (or more correctly duplicating) about a dozen mains leads in the first two cabinets, where the rubber insulation has perished. We have adopted a policy of replacing with white plastic sheathed cable, clearly distinctive from the original cable, and leaving the latter in place but not connected. Once these replacements are made, further progress with establishing power to the system can proceed.

Small-Scale Experimental Machine

Chris Burton

At the request of the Museum of Science and Industry, we have committed to giving live public demonstrations of the replica computing machine at 1200 and 1400 every Tuesday. A rota of team members has been set up to do this, to distribute the burden. Longer term, a start has been made on training a group of about eight volunteers to give these demonstrations, and where skills are appropriate to perform running repairs on the machine.

Training includes health and safety, facing the customers, switching the machine on and off and running demonstration programs. Further training on the way the machine works and hints about fault-finding will be provided to those with the aptitude. We are always pleased to hear from additional volunteers who would like to become authorised demonstrators.

Through this regime of weekly running of the machine, we have been pleased at the reliability. A few faults have occurred, but generally they are identifiable and repairable. Surprisingly, the main cathode ray tube

store has improved over the past few weeks, though no work has been done on it.

The remote control interface which has been under development by Dr Alan Knowles has been connected to the machine. This duplicates the operation of all the control switches and push-buttons under control of a PC program. The system permits automatic loading and operation of programs for realistic demonstrations.

Pegasus Working Party

Derek Milledge

After completion of package testing and adjustment, reported in *Resurrection* issue 20, we found that Pegasus was able to run successfully 'on margins' at reduced voltage, showing that the machine is now working better than it ever was in the Old Canteen.

With the drum working well we were able to install the 1963 Engineers' Test programs on the isolated drum store in addition to the 1957 tests, and run through both successfully. Unfortunately our last meeting ended with a recurrence of drum failures at certain familiar locations, so we still have to sort that out.

With help from the Documentation Department at the Science Museum we now have all available engineering diagrams at Blythe Road, though they do not include the modifications made when our Pegasus was used for factory tests in the early 1960s.

We have started to look at the character handling instructions which were additional to the original Pegasus 1 specification. Simple tests now work, but failures on special cases have yet to be resolved.

We had expected Pegasus to go into the new 'Making of the Modern World' gallery at the Science Museum in the year 2000, but this is now uncertain. A site in the gallery has been offered, but it is too small to maintain and run the machine. Negotiations are continuing with the aim of obtaining sufficient space to show a working Pegasus.

Software Conservation

Dave Holdsworth

The Society is contemplating a new venture, designed to reincarnate the genie in the software of yesteryear. Put more directly, the intention is to

use the techniques of long-term data preservation to keep the software in a machine-readable form indefinitely, and to implement software emulation so that future generations can experience George 3 or try their hands at Mercury autocode.

Effort for emulation becomes available in 2001, but we need to be preserving data *now*. We appeal for machine-readable vintage software now, such as operating systems and compilers. As far as delivery media are concerned, half-inch nine track tape should pose no problems, and seven track is probably accessible also. Real gems (such as Atlas) on other media may be worth the effort of trying to get old drives to work.

Will anyone who can help please contact me by telephone (0113 233 5402) or e-mail (D.Holdsworth@leeds.ac.uk).

“Another ICL Anthology”

Nicholas Enticknap

Our indefatigable Secretary has produced a second compendium of anecdotes about the lighter side of working for ICL and its antecedent companies. Hamish Carmichael tells us, among other things, how easy it was to sell Powers Samas equipment to Robert Maxwell; how Powers Samas finally broke into the US; why ICT did not take over the computer side of Ferranti till 1963; why Leo was not programmed to cater for the millennium bug; who the managing director’s daughter married; how Arthur Humphreys was disappointed by the egg boiling machine; how ICT influenced the screenplay of the film “2001”; how the George operating system got its name; why documentation is like sex; and why the Danish software company MD did not want milk and sugar.

If this sort of thing is your kind of history, you can obtain “Another ICL Anthology” from Laidlaw Hicks Publishers at 63 Collingwood Avenue, Tolworth, Surbiton, Surrey KT5 9PU. Cheques should be made payable to JWS Carmichael, for £12.50 (UK), £13.50 (elsewhere in Europe) or £16.50 (anywhere else). The prices include postage and packing.

Bombe Rebuild Project — Appeal for More Help

The Bombe Rebuild Project is making a further appeal for help from anyone with skills in the following areas.

Adviser in how best to get ‘plastic’ mouldings manufactured. There are over 100 of each item required, so tooling and/or moulds are anticipated. 3D drawings are available but we need help and advice with the next step.

Project Facilitators/Estimators are needed now to take copies of mechanical drawings out to potential manufacturers, to obtain quotations and, when prices are agreed and funds are available, progress the manufactured parts through to delivery. The work will be split down into manufacturing or material types, such as gear cutting, precision casting, turned parts, sheet metal, insulating material.

Fund Raiser to seek out prospective sponsors and co-ordinate presentations and submissions, and thereby raise the funds necessary to complete the rebuild, together with some additional funds for ongoing maintenance and repair. It is emphasised that we will have to approach relatively large organisations in order to raise an amount that is expected to run well into six figures.

Small specialist part manufacturers with access to a modern, precision, well equipped workshop.

Small ‘simple’ part manufacturers There are many small, simple steel and insulating material parts that need to be made. These can often be made in a small domestic workshop as they only entail cutting and drilling. However, accuracy is important. Material and drawings will be supplied.

Anyone who can help with any of these requirements should contact John Harper by e-mail at bombe@jharper.demon.co.uk, by phone on 01462 451970 or by fax on 0870 055 4870.

<p style="text-align: center;">Editorial contact details</p>

<p>Readers wishing to contact the Editor may do so by fax to 0181-715 0484 or by e-mail to NEnticknap@compuserve.com.</p>
--

Forthcoming Events

9-10 January 1999, and fortnightly thereafter Guided tours and exhibition at Bletchley Park, price £3.00, or £2.00 for concessions

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

26 January 1999 North West Group meeting on “The Leo Computer”

Speakers include John Aris and George Manley

20 February 1999 Newcomen Society meeting on “Evolution of Modern Electronics”

23 February 1999 North West Group meeting on “Computers on Display”

A behind the scenes look at developing an exhibition

26 February 1999 London meeting on “The Distributed Array Processor”

Speaker Dennis Parkinson: readers should note that this meeting takes place on a Friday rather than the usual Thursday

23 March 1999 North West Group meeting on “The Transputer”

Speaker Iann Barron

15-16 April 1999 EDSAC99

See page 22 for further details

The North West Group meetings will take place in the Conference room at the Manchester Museum of Science and Industry, starting at 1730; tea is served from 1700.

The Newcomen Society meeting on 20 February will take place at Imperial College, London, from 1030 to 1730.

For more information on London meetings, readers should refer to the insert enclosed with this issue.

Queries about London meetings should be addressed to George Davis on 0181 681 7784, and about Manchester meetings to William Gunn on 01663 764997.

FTP, Web and E-mail Addresses

The Society has its own World Wide Web (WWW) site: it is located at <http://www.cs.man.ac.uk/CCS/>. This is in addition to the FTP site at <ftp.cs.man.ac.uk/pub/CCS-Archive>. Our Web site includes information about the SSEM project as well as selected papers from *Resurrection*. Readers can download files, including the current and all past issues of *Resurrection* and simulators for historic machines.

Readers of *Resurrection* who wish to contact committee members via electronic mail may do so using the following addresses.

Chris Burton: chris@envex.demon.co.uk

Martin Campbell-Kelly: mck@dcs.warwick.ac.uk

Hamish Carmichael: hamish.carmichael@bcs.org.uk

George Davis: georgedavis@bcs.org.uk

Nicholas Enticknap: NEnticknap@compuserve.com

John Harper: bombe@jharper.demon.co.uk

Dan Hayton: Daniel@newcomen.demon.co.uk

Len Hewitt: leonard.hewitt@virgin.net.

Dave Holdsworth: D.Holdsworth@leeds.ac.uk

Roger Johnson: r.johnson@bcs.org.uk

Adrian Johnstone: adrian@dcs.rhbnc.ac.uk

Simon Lavington: lavis@essex.ac.uk

Brian Oakley: brian.oakley@ukonline.co.uk

Tony Sale: t.sale@qufaro.demon.co.uk

Robin Shirley: r.shirley@surrey.ac.uk

John Sinclair: john.eurocom@dial.pipex.com

John Southall: jsouthall@bcs.org.uk

Doron Swade: d.swade@ic.ac.uk

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

Editor – Nicholas Enticknap

Typesetting – Nicholas Enticknap

Typesetting design – Adrian Johnstone

Cover design – Tony Sale

Printed by the British Computer Society

©Computer Conservation Society

Committee of the Society

Chairman **Brian Oakley CBE FBCS**, 120 Reigate Road, Ewell, Epsom, Surrey KT17 3BX. Tel: 0181 393 4096.

Vice-Chairman **Tony Sale FBCS**, 15 Northampton Road, Bromham, Beds MK43 8QB. Tel: 01234 822788.

Secretary **Hamish Carmichael FBCS**, 63 Collingwood Avenue, Tolworth, Surbiton, Surrey KT5 9PU. Tel: 0181 337 3176.

Treasurer **Dan Hayton**, 31 The High Street, Farnborough Village, Orpington, Kent BR6 7BQ. Tel: 01689 852186.

Science Museum representative **Doron Swade CEng MBCS**, Curator of Computing, The Science Museum, Exhibition Road, London SW7 2DD. Tel: 0171-938 8106.

Chairman, Elliott 803 Working Party **John Sinclair**, 9 Plummers Lane, Haynes, Bedford MK45 3PL. Tel: 01234 381 403.

Chairman, Elliott 401 Working Party **Chris Burton CEng FIEE FBCS**, Wern Ddu Fach, Llansilin, Oswestry, Shropshire SY10 9BN. Tel: 01691 791274.

Acting Chairman, Pegasus Working Party **Len Hewitt MBCS**, 5 Birch Grove, Kingswood, Surrey KT20 6QU. Tel: 01737 832355.

Chairman, DEC Working Party **Dr Adrian Johnstone CEng MIEE MBCS**, Royal Holloway and Bedford New College, Egham, Surrey TW20 0EX. Tel: 01784 443425.

Chairman, S100 bus Working Party **Robin Shirley**, 41 Guildford Park Avenue, Guildford, Surrey GU2 5NL. Tel: 01483 565220.

Chairman, Turing Bombe Project **John Harper CEng MIEE MBCS**, 7 Cedar Avenue, Ickleford, Hitchin, Herts SG5 3XU. Tel: 01462 451970.

Chairman, North West Group **Professor Frank Sumner FBCS**, Department of Computer Science, University of Manchester, M13 9PL. Tel: 0161 275 6196.

Meetings Secretary **George Davis CEng FBCS**, 4 Digby Place, Croydon CR0 5QR. Tel: 0181-681 7784.

Editor, Resurrection **Nicholas Enticknap**, 4 Thornton Court, Grand Drive, Raynes Park SW20 9HJ. Tel: 0181-540 5952. Fax: 0181-715 0484.

Archivist **Harold Gearing FBCS**, 14 Craft Way, Steeple Morden, Royston, Herts SG8 0PF. Tel: 01763 852567.

Dr Martin Campbell-Kelly, Department of Computer Science, University of Warwick, Coventry CV4 7AL. Tel: 01203 523196.

Professor Sandy Douglas CBE FBCS, 9 Woodside Avenue, Walton-on-Thames, Surrey KT12 5LQ. Tel: 01932 224923.

Dr Dave Holdsworth MBCS CEng, University Computing Service, University of Leeds, Leeds LS2 9JT. Tel: 0113 233 5402.

Dr Roger Johnson FBCS, 9 Stanhope Way, Riverhead, Sevenoaks, Kent TN13 2DZ. Tel: 0171-631 6709.

Professor Simon Lavington FBCS FIEE CEng, Department of Computer Science, University of Essex, Colchester CO4 3SQ. Tel: 01206 872677.

Graham Morris FBCS, 43 Pewley Hill, Guildford GU1 3SW. Tel: 01483 566933.

John Southall FBCS, 8 Nursery Gardens, Purley-on-Thames, Reading RG8 8AS. Tel: 0118 984 2259.

Ewart Willey FBCS, 4 Sebastian Avenue, Shenfield, Brentwood, Essex CM15 8PN. Tel: 01277 210127.