


Issue Number 40

Summer 2007

Computer

RESURRECTION

The Bulletin of the Computer Conservation Society



science
museum

 **BCS**
THE BRITISH COMPUTER SOCIETY

THE MUSEUM
OF SCIENCE &
INDUSTRY
MANCHESTER

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 voluntary subscriptions from members, a grant from the BCS, fees from corporate membership, donations, and by the free use of Science Museum facilities. Some charges may be made for publications and attendance at seminars and conferences.

There are a number of active Working Parties on specific computer restorations and early computer technologies and software. Younger people are especially encouraged to take part in order to achieve skills transfer.

Resurrection

The Bulletin of the Computer Conservation Society

ISSN 0958-7403

Number 40

Summer 2007

Contents

Editorial	2
<i>Nicholas Enticknap</i>	
News Round-Up	3
Society Activity	5
CCS Web Site Information	7
Production Control at Letchworth	8
<i>Hamish Carmichael</i>	
June – a Home-Grown Operating System	21
<i>Andrew Colin</i>	
Memories of the Elliott 803B	25
<i>Graham Phillips</i>	
Letters to the Editor	28
Forthcoming Events	31

Editorial

Nicholas Enticknap

Welcome to issue 40 of *Resurrection*. We concentrate in this issue on the days when computing was just beginning to become established in British industry. A decade had passed since the first tentative steps were taken towards the modern IT world at Manchester and Cambridge universities, and the boardroom at Lyons witnessed the first glimmerings of appreciation of what the computer might be able to do for business.

This period, the late fifties and early sixties, was a time when senior corporate management was beginning to understand how computers might transform industry, but the technology was still as crude as that of the Edwardian motor car and computer professionals required all their dedication and ingenuity to keep their systems running. Both these themes are explored in the following pages.

Our first article describes a production control application that was conceived as far back as 1956. This is a story of a computer manufacturer, ICT, practising what it preached and applying its own innovative technology to its own problems. Hamish Carmichael tells the story using the company's own words, providing a fascinating insight into the thinking of the time.

This article formed the basis of the talk to the Society last November: the other two are contributions from readers recounting their own experiences at the coal face in the sixties.

Former university lecturer Andrew Colin describes how, faced with a primitive early operating system unsuited to the demands of an academic environment, he and his colleagues rolled up their sleeves and developed their own tailor-made alternative, which reduced the minimum job time from three and a half minutes to 12 seconds.

Graham Phillips talks from the very different perspective of an Elliott Brothers engineer and tells of problems and excitements involved in keeping Elliott 803s going at a time when a trip up the M1 was still a memorable experience.

These stories serve as a reminder of how far we have come in half a century. We shall be celebrating that progress in the big IT event of this year, the BCS Golden Jubilee conference, which is now almost upon us. We trust that all readers who are planning to go to either the Bletchley Park part or the London part of BCS@50 enjoy what is sure to be a memorable occasion.

News Round-Up

The replica Colossus will play a major role in the new National Museum of Computing at Bletchley Park, thanks to a generous donation of £75,000 from the British Computer Society. The money has been used to support the re-installation of the world's first digital computer in its original location in Block H at Bletchley.

101010101

This is the culmination of 13 years work by CCS Vice-Chairman Tony Sale and his team of volunteers, adding up to 6000 man-days of effort. The replica Colossus is now situated in the location occupied during the War by Colossus number 9.

101010101

There have been further changes to ticket prices at Bletchley. Tickets now allow free repeated returns during the 12 months following the date of purchase. Also the age limit for children to be free is now increased to 12, and the price for those from 12 to 16 is reduced to £6 (other concessions are still £8). Accordingly the family ticket has been reduced to £22.50. On site parking is £3 each visit. There is no extra charge for special events in the day, but evening events such as lectures, music and fireworks night do have a charge.

101010101

HRH Duke of Kent will officially switch on the reconstructed Bombe on Tuesday 17 July, one week after the BCS Golden Jubilee Conference.

101010101

The BCS 50th anniversary conference will take place from Thursday 12 July to Saturday 14 July 2007. The first day's proceedings of BCS@50 will be held at Bletchley Park, with the next two days taking place at the BCS offices in London. Admission is £28 for each segment.

101010101

We regret to report that Derek Eldridge died on 2 April. Derek worked for IBM, Ferranti and ICT in the 1950s and 1960s, and contributed an article to *Resurrection* 16 about his role in the design of the ICT 1900 series.

101010101

David Link of Cologne has built an emulator for the Ferranti/Manchester Mark I, and has proved it by running some original programs for the machine, including Christopher Strachey's love letters program. We hope to persuade Dr Link to give a talk to the Society on his work.

101010101

The Real Time Club will be holding a celebration on Tuesday 26 June "to celebrate 40 years of iconoclastic dining". The event takes place at The National Liberal Club in London. Prices are £85 for one person, or £140 for two.

101010101

John Backus, a pre-eminent figure in the history of programming, died on 17 March at the age of 82. Backus was best known for his work as leader of the team that developed the first Fortran compiler, and as co-developer with Peter Naur of the Backus-Naur programming notation.

101010101

Society Activity

Pegasus Working Party

Len Hewitt and Peter Holland

Pegasus has been run as usual at the museum at fortnightly intervals. There have been few problems, mostly of a minor nature. Recently we have had trouble with the output paper tape punch. This has been replaced with a spare punch from our store; the original has still to be repaired. The motor generator has been serviced and the Star/Delta changeover cut-out switch has been replaced. This had been giving us trouble on start-up for some time, but replacement with a higher capacity switch seems to have cured the problem. We plan to continue to display the machine to an increasing number of visitors brought to us by the better publicity we now receive from the museum. We thank them for this and for their assistance with the motor generator.

Contact Len Hewitt at leonard.hewitt@ntlworld.com .

Bombe Rebuild Project

John Harper

Almost all of our effort now is going into preparations for the BCS@50 event this coming July. We are now able to provide a presentable demonstration of the machine running a single menu, and we have plenty of time still available to elaborate on this. Running more complex menus would be of great interest to the team but would not improve on the simple demonstration to take place in the time that visitors would be able to spend with us.

On the technical front much of our time recently has been spent getting the settings of what are called the Stop and Cancel circuits correct. The photo overleaf shows the mechanical part of this subsystem. With it are 26 relays. We have found that the brushes that sweep the commutator surface have to be very precisely set and be of the correct end profile in order to achieve the correct operation of the clutch and other circuits. Timings have to be within about two milliseconds for correct operation, which is a very tall order for electromechanical equipment.

We initially used recovered contacts in the associated 26 ‘stop’ relays, but these were not reliable so we had to make new parts. For these to be correct as in the original drawing we had to find a method of welding tungsten tips onto steel backs. These tips were only 0.04" thick and 0.1" diameter, so were quite fiddly. We succeeded by making special electrodes for a car spot welder. After cleaning up and polishing the final items we then riveted them on to new relay spring leaves. We required 104 contacts, so this took a lot of team effort. This is now behind us and full system testing is currently taking place, bringing into play parts of the machine that have not yet been fully exercised.




Figure 1. Stop and Cancel Unit

For the July events we still hope to move to a larger area. This is mainly to allow more people to gather around the machine at one time, but it will also give us more space for other equipment and displays. However as I write this move is not definite. Problems with an old BP building currently being refurbished being in a worse state than originally thought have slowed down things moving around the BP site.

As I said at the beginning of this report, our main effort will now be on displays and presentations. Our fully completed and operational Checking Machine will be part of this display. This might well be used to demonstrate how a Bombe ‘stop’ can be checked to see if it were valid and therefore worthy of being submitted as a potential setting that would allow German traffic to be decoded.

We will also demonstrate a Typex machine that we modified to work as a German Enigma machine. This will decode German messages using the settings verified by the Checking Machine. Other display panels etc. will complete the whole display arrangement.

Our Web site is at www.jharper.demon.co.uk/bombe1.htm.

CCS Web Site Information

The Society has its own Web site, located at www.bcs.org.uk/sg/ccs. It contains electronic copies of all past issues of *Resurrection*, in both HTML and PDF formats, which can be downloaded for printing. We also have an FTP site at [ftp.cs.man.ac.uk/pub/CCS-Archive](ftp://cs.man.ac.uk/pub/CCS-Archive), where there is other material for downloading including simulators for historic machines. Please note that this latter URL is case-sensitive.

Production Control at Letchworth

Hamish Carmichael

This contemporary description of an early computerised production control application illustrates some of the typical problems faced in the pioneering days and the lateral thinking involved in solving them. It presents an interesting snapshot of the time when the potential of the computer was just beginning to be appreciated by industry.

Many readers will recall an article in *Resurrection* 36 by Hugh McGregor Ross in which he started by saying that when computers were invented nobody really knew what to do with them.

I've been pondering that remark off and on, and I think I've come up with the proper response. It's a quotation from Kipling, who gives the great Pyecroft the line:

"I wouldn't exactly call him a liar, but I might say that he and the Admiralty are pretty alike in their statements".

The point is that Hugh was a Ferranti man, and Ferranti was never really in the data processing game.

However, from 1950 onwards an increasing number of punched card installations included not only the conventional mix of sorters, tabulators, collators, reproducers and interpreters, but also the early electronic calculators. Following my lifetime's principle I shall ignore IBM machines: I'm thinking of the Hollerith 540, 541, 542, 550 and 555 and the Powers Samas EMP. There must have been somewhere about a thousand of these installed in the UK alone.

And forward-thinking data processing managers knew both that they wanted more and also that they wanted something different. Anyone who bears the mental scars after plugging up a 555 for a tricky job will remember what a technological dead end that machine was.

So while few managements were as far-sighted and clear-headed as Joe Lyons, and while most boards of directors shared the typical British attitude of not wanting to be the first patient to try out a new treatment at

the dentist, there were plenty of data processing managers pushing to adopt the new computing technology.

The tables of early installations show that the number of data processing machines quickly overtook and then far outran the number of computers used for what a Ferranti man would call real computing.

I want to discuss one such data processing application. I found a write-up of it in the ICL archive. It probably dates from 1959. I present a condensed version of it more or less straight, because I think it's interesting to see not only what the application did and how, but also to notice what the British Tab people of that time thought they needed to explain and justify to their customers.

Background

For over half a century the former British Tabulating Machine Co Ltd (now joined with the former Powers-Samas Accounting Machines Ltd as International Computers and Tabulators Limited) had been manufacturing punched card equipment. More recently it had also established itself in the field of electronic data processing – that of Calculators and Computers.

The growth of the Company since the last war had been phenomenal, and still it goes on in the combined Company. This rapid and continuous expansion provided and still provides many problems for Production Management.

In 1956 the Directors decided that an order should be placed for a Type 1201 computer for use in the organising and communicating of information about the productive activities of its engineering factories.

The objects were:

- *to cope with the rapid growth in demand for the Company's products and the consequential expansion of manufacturing capacity to meet such demand;*
- *to exercise closer control on inventory values (it is all too easy to lock up vast sums of money in inventory as an insurance against increases in future demand); and*
- *to deploy the capital resources of the Company – money, machines and facilities – to the greatest advantage in order to obtain the optimum yield.*

The Manufacturing Unit

For the purpose of the application, a number of the Company's engineering factories comprise together a single manufacturing unit to the capacity of which the requirements of the consumer units – Sales, Field Engineering, Technical – are to be related. Nevertheless, these factories are separate and distinct producers:

- *No 1 Factory at Letchworth, Herts, where production is confined to tabulators in their various types.*
- *No 8 Factory at Castlereagh near Belfast in Northern Ireland, the home of the ancillary machines – Sorters, Collators, Summary Punches, Reproducers and Interpreters. Here also are produced Electronic Multipliers and Calculators, whilst a new production line for computers of the 1200 series is being laid down to supplement and ultimately to supersede the existing one at our electronics factory in Stevenage.*
- *Factory 1b for the manufacture of components, a tributary factory that feeds factories 1 and 8 with items for embodiment in the actual machines they build.*

The scheme is planned to embrace production at all three factories, but 1b factory is the guinea pig for testing its practical working.

Approximately 600 different end products are made in this factory, which is also at Letchworth. Half of them are cable assemblies consisting largely of plastic covered wires, string and connecting tags.

Programme Breakdown

Our starting point is an agreed consolidated programme setting out the requirements for a 12 month period, by Sales, Field Engineering, and Technical sources, of specified end products.

Of these requirements those for the first four months are to be treated as positive and firm, those for the subsequent eight months as tentative and subject to revision. The overall programme is reviewed and revised thrice yearly and upon each revision the positive requirement is brought forward four months in time.

The procedure is to punch up a card for each end product in the programme, identifying it and stating the number of units thereof called for in each of a pre-determined number of periods of time ahead. In this form the information can be fed into the Computer, stored in specific identifiable locations on the track of the magnetic drum and held there to serve as a multiplier in subsequent calculations.

So much for the multipliers; now to find the multiplicands. Records are available of the unit content of each type of end product, ie the number or quantity of each item – assembly, piece part, or material – required to make one-off of that end product. Each such component item is identified by a part number.

Subject to detail modifications from time to time arising from specification changes or from the fact that tabulators are tailor-made to individual specifications, this record of the unit content of the individual end-product is more or less permanent, is a constant in fact.

The record in this form is unsuited to our purpose since we are not concerned with end products but with establishing what things have got to be procured to make them. We have therefore twisted the data round to produce a different form of record showing under item – identified by part number – all the different end products in which that item is used, with the number-off of the item for each end product.

Having got this constant information into a more useful form we then punch up cards to form a permanent file of master cards which, subject to detail modification on a specification change, can be used and re-used as often as required for a new breakdown. First a pack of conventional cards is produced, with each card containing part number, end product number, and quantity-off per end product. But because of the numbers of parts and of end products, this pack is too large to be conveniently stored between successive programme breakdowns. It is therefore input to a preliminary program in which the 1201 itself summarises the information, and produces what are called Binary Part Master Cards.

Since the design of the Binary Part Master Card is somewhat novel and because of its importance in the computer application, it should be considered in some detail.

Two forms of punching are used:

- 1. decimal – or conventional – punching where the meaning of a hole is determined by its position in both a particular column and a particular line;*
- 2. binary punching, where a binary number can be punched in a single line, with a hole for each '1' and absence of hole for each '0'.*

The Part Number, Factory Location and some other fields use conventional decimal punching. The Quantity field (in columns 11 - 50) and the Drum Location field (in columns 61 - 70) use binary punching,

and because each binary number only uses one line of the card, the 12 lines of the card can be used to hold 12 binary numbers. Because in practice no quantity in this programme breakdown can be larger than 10 binary digits, the 40 columns of each quantity line can be used to hold four binary quantities. The 12 lines of the card thus mean that a single card can hold 48 quantity fields. The 40 bits of each quantity line conveniently correspond to the word length of the 1201.

The Drum Location line contains the address on the drum where the corresponding quantities are to be stored.

The quantity sub-sections or boxes, and the corresponding drum locations, automatically identify positionally different end products. So position 01 always represents one specific end product, position 02 another, and so on. Thus by location alone we can represent 48 end products on a single Binary Part Master Card.

Because we have to be able to handle up to 600 end products, there can be more than one Binary Part Master Card for each part number. An extremely common part could require 13 Binary Part Master Cards.

In our file of Binary Part Master Cards we have a permanent record for each item identified by a Part Number in our Parts List, showing which end products it is used in, and how many of it are used in each end product. The rest is simple. We feed this information into the Computer, and relate it to the information already stored there telling us the quantity required for each end product. By successively multiplying together the paired quantities – end product requirement quantity times number-off of constituent item per unit of end product – and accumulating the results of the successive multiplications, we get the total requirement for each item.

Thus we establish, by Part Number, our overall requirement of physical items to enable us to make end products to the consolidated Programme of Requirements. If we were dealing with a factory starting up from scratch, the gross requirement so established would be the actual requirement. But we are making our demand upon a manufacturing unit which is already an active going concern, with Work in Progress, orders outstanding on Suppliers, and Stores holdings.

Not all of the items from these three sources will be absorbed in current work. Therefore, having now established the gross requirement by Part number, we subtract from such total the free stocks of the item available to us, so obtaining for every Part Number the true net requirement. This is the number or quantity which serves as the basis for

initiating procurement action, either by the purchase department or by production.

The method here outlined, which gives a one-shot breakdown by directly relating the constituent item to the quantity of end products for which it is required, is novel and does demonstrate how re-thinking a problem in order to put it on the Computer can contribute to more effective results. For us a Programme Breakdown for the whole manufacturing unit which would by other methods take several weeks will now be carried out in a number of days.

Plant Loading

We take the output from the Programme Breakdown, separate out those items to be bought in from suppliers, and are left with a statement of requirement of made-in items in the form of a file of punched cards taken from the gang punch unit of the Computer. This file is the starting point of the Plant Loading procedure. It tells us what is to be made, how many are to be made, and where they are to be made, but we do not yet know in what order they are to be made or when to start making them.

In the factory no one is concerned with the items to be made; they are interested solely in the manufacturing operations. Given that these operations are performed correctly and in the correct sequence, the inevitable consequence is the Part. It is therefore necessary to establish what the known requirements represent in terms of operations.

There is one point about loading which should be made for those who are familiar with computers but not so conversant with what happens on the shop floor. The quantities to be loaded of any given item do not necessarily bear an exact relationship to the requirements determined by the Programme Breakdown procedure.

The preparatory work leading to the performance of an operation takes time and costs money. It can take several hours to set up a machine tool to perform a particular operation in a few minutes. In such cases it is not worthwhile setting up – and breaking down, for that too takes time and costs money – every time a one-off or a few-off of a part is required.

On the shop floor, once a machine is set up, the urge is to use it to churn out as many of the items as possible. This urge has to be resisted, since it can result in overloading the Stores, tying up money in inventories, and obsolescence loss. There was a recent case of a 20 years supply of a particular small pressing rendered obsolete after six months by a change in design.

However, experience and close study have established for most items an economic manufacturing quantity or economic batch for which it is worthwhile to set up a machine but which is not too remotely related to usage. Obviously, enough must be made to cover the net requirement established by the Programme Breakdown. If the economic batch is more than the net requirement then the balance will be held in Stores for absorption in a subsequent requirement. If the batch is less, then two or more batches will be loaded and again any surplus will be held in Stores.

So in the loading we are concerned with the time taken to perform any particular operation on an economic batch. For each item, piece part or assembly, it is necessary to record on a Binary Operations Master Card a correct sequence of operations, defining for each operation:

- the set-up time in hours to the nearest tenth;*
- the allowed time for one-off; and*
- the time interval between a particular operation and the next in sequence; this allows for handling between machines and similar not precisely definable events.*

The sum of these three times for a particular operation is the overall operation time, and the sum of all the overall operation times should be the time required to make an economic batch.

We now determine when in the overall production cycle each part or assembly has to be made. We start with the agreed delivery date, when the end product of Factory 1b must be in the hands of Factory 1 or 8 to enable them to complete their tasks. Then we work back, assembly by assembly and part by part. This analysis carried to its logical conclusion establishes on the computer the time at which any part or assembly must exist in order for it to be available for use in the next production stage.

The Requirements Cards and the Binary Operations Master Cards together provide the means of establishing a picture in absolute time of how the requirements should be met and when, in relation to some production cycle. This presupposes a hypothetical factory standing by completely equipped and fully available to cope with any demands we care to make upon it. But we have in Factory 1b an actual working manufacturing plant which is already engaged upon making end products.

To take account of these conditions we have a third file of cards, the Machine Tool Capacity Cards. These record the overall capacity, and the existing load thereupon, in terms of machine tool groups. The

difference between overall capacity and existing load represents the time available on machine groups during given weeks of the production cycle.

So with three files of cards now available:

- 1. Requirements from the Programme Breakdown*
- 2. Operations in correct sequence to make each part or assembly*
- 3. Capacity of groups of machines each week*

we now use the Computer to calculate the overall load – existing load plus new load – for each machine tool group during each week of the production cycle.

This first calculation enables us to detect at once if the added requirement will result in overload anywhere along the line and, if so, where. If there is an overload, we can decide at this early stage how to deal with it.

Having established the broad picture, detailed plant loading is effected at more frequent intervals. Starting at any given week of the production cycle, and taking advantage of the discriminatory facilities offered by the Computer to test load requirements against available capacity, one works backwards week by week until the total load is successfully absorbed.

The output of the Computer at this Plant Loading stage comprises:

- A Load Statement – for the guidance of Production Management, not for issue to the shops – showing the load on different shops, sections and machines week by week over the production period.*
- A Statement of Requirements for jigs, tools, fixtures and other facilities for the carrying out of each operation – in other words a subsidiary load statement.*
- Operation Job Cards – for each operation on each part or assembly. These are not produced for the whole of the production period; only one or two weeks' cards are needed at a time.*

Progress Control

Any resemblance between the load statement produced by the computer and what actually happens on the shop floor depends on:

- *precise definition of what is to be done and when – the operation job card and its accompanying documents; and*
- *effective progressing of the work to see that it is actually carried out as prescribed.*

The load statement and operations job cards together set out and define the long term strategic purpose of the exercise, covering the whole of a production cycle – or perhaps even more. These documents are subject to modification in the light of events, eg changes due to a crash programme, or a change in specification, or perhaps merely the re-introduction of items of which a shortfall has been reported. This strategic purpose is not communicated to the shop floor.

The short term tactical purpose of the exercise, the defining to the shops in precise terms what they are to do next week, how and in what order, is achieved by issuing to them week by week:

- *that section of the load statement covering all jobs to be done next week. This is issued one week before the work is actually to be done.*
- *operation job cards covering all jobs to be done next week. These are issued one day before the week in which the work is actually to be done.*
- *Drawings, Layouts, Batch Label, Material Requisitions covering all aspects of the job for which the Operations Job Cards are issued. The first three will accompany the job throughout its whole course from start to entry into store.*

You can feed a Charge Hand or Foreman with documents setting out precisely what is to be done, when and how, but you do not assume that he will turn out the work as required. It is necessary to monitor the happenings on the shop floor and report these back to Production Management in the form of feed back called Progress Control.

The Operations Job Card is the key to the system of Progress Control. Four copies of this are produced, as interpreted punched cards, in the computer room by gang-punching from the original, which forms the punched card output of the Computer at the Plant Loading stage. Of these copies

- a) is retained in the Computer Room*
- b) is retained in the Central Control Room*
- c) is issued to the Timekeeper*

d) is the shop card copy, issued to the foreman with the necessary drawings, batch label, etc and constituting his authority to do the work.

When the first individual part of a batch has been machined it is submitted for inspection. If he passes it, the Inspector initials the Shop Card as authority to proceed with the making of the batch. When the batch is completed, it is passed in for inspection and counting, and the results are entered by the Inspector on the Shop Card. The job with its associated documents is then available for the next operation.

The Shop Card is returned to the central Control Room where the quantity passed good is noted and marked (for mark sensing) on card copy (b). The Shop Card is then passed to the Timekeeper who marks quantity good on his own copy (c) and sends this to accounts for costing and other purposes. The Shop Card he files and retains for one month in case queries arise as to bonuses. It can then be destroyed.

The marked copy (b) is passed to the computer room where the marking is sensed and punched into the card. The computer room now holds its own copy (a) of all the Operations Job Cards issued to the shops, and the control room copy (b) of all cards returned from the shops. A simple collation of (a) against (b) with the throwing out of unpaired cards, ie those representing jobs not yet completed, shows what jobs are behind programme. These are then listed, section by section, in priority order, and each section is issued with a precise statement of any jobs in that section that are behind schedule.

If conditions in any section become so serious that the work of other sections is likely to be held up, then re-loading may be necessary. By manual or conventional mechanised methods this might be a major operation; with the Computer it presents little difficulty. Since the shops are only concerned with the current and the following week's work, reloading is largely a matter of re-assessment within the control room and the computer room. The shops need not even become aware of it.

The Shop Card copy (d) for the last operation serves to book the finished parts into store by quantity. It is in effect the Goods Made Card.

This method of Progress Control is simple and straightforward, but the method is only possible given the ability at any time quickly to review and re-assess the whole strategic long-term picture and to prepare operations job cards to meet any new situation. This ability only the Computer provides.

Stock Control by Computer

Our outlook and approach, coloured by our experience of using the Type 1201 computer, does differ from and does represent an advance upon conventional routine stock control procedures.

This outlook and this approach are based upon certain premises:

- a) That post mortems may be interesting but are seldom useful. One cannot exercise control over what has already happened. The object is, so far as possible, to control what is happening now and what is to happen in the future. For this purpose the conventional stock record, which is a historical record of past events, is of very limited use.*
- b) That “Dog bites man” isn’t news; “Man bites dog” is. Only departure from an accepted norm is of interest. No useful purpose is served by investigating any stock item unless some action has to be taken.*

From these premises we have endeavoured, with the aid of the Computer, to develop an efficient system of Stock Control, one which will effect more frequent adjustment of stock records and provide speedy and effective control of stocks by:

- projecting our investigations into the future;*
- taking account of more variables than is customary; and*
- indicating specifically only those items upon which some action must be taken, and what action.*

Such an approach is more comprehensive and more realistic than is usual in these applications.

The system we are developing provides for:

- 1) a weekly review of the position relative to stocks of raw materials, of all other bought-out items, and of all long-term sub-contracts. This review covers, by period, 10 periods representing in all requirements for some 21 months ahead.*
- 2) a daily review of the position relative to stocks of all made-in items (piece parts or assemblies), such review covering by period 10 weekly periods.*

The object of the exercise is to make weekly or daily adjustments to stock records so that, by making use of the Computer’s discriminatory powers, decisions are automatically taken as to dealing with any variations from programmed stocks, either by absorption into the next

order or immediate initiation of a new order, in either event providing a factual basis for action.

There is nothing exceptional about the actual procedure as at present envisaged. The Stock Balance Card is just another of these part-decimal part-binary cards, two of which have already been discussed.

It is not suggested that we have exhausted all the potentialities of the Computer as an aid to Stock Control. Research into the subject still continues. The procedure outlined does, however, appear to represent a considerable advance upon anything hitherto attempted save in very simple cases.

Summary

In the foregoing description of the application of the ICT Type 1201 computer to Production Control in one of our own engineering factories we have attempted to cover what seemed to us the important aspects, and to deal with the major problems of the subject. The solutions we propose are those which, in the light of our present knowledge, would appear most practicable and workable.

In conclusion we would offer for your consideration the reason why, in our opinion, the Computer can be, and in the future is likely to be, so potent a tool of management. That reason may be summed up in one word: freedom.

Properly used the Computer can afford Management:

- *freedom from the time barrier interposed by conventional methods of organising and communicating information between the facts and those who make use of them;*
- *freedom from the irksome task of routine decision making. Where decisions arise naturally out of the situation they can, automatically, be taken by the machine, and be correctly taken on the basis of facts, not hunches;*
- *freedom from the vast volume of paper work, relating for the most part to past events about which nothing useful can be done, or to present routine events about which information is not required anyway.*

In buying or renting a computer one is buying time, that precious commodity of which management is so frequently so short. The ever-increasing tempo of modern industry is such that only those managements prepared to adopt modern methods and techniques can hope to compete

and survive. One such aid to survival is that which has formed the subject matter of this presentation – the ICT Type 1201 Computer.

Postlude

Looking back at what ICT said in 1959, I find it remarkable that there was so little technical detail and virtually no mention of the actual computing technology, which we would expect to be at the forefront of everyone's attention. There were a few bits of clever stuff in the design of particular punched cards and one or two casual references to drum addresses, but that's it. One can deduce that the aim was to show how the arrival of the computer was not going to mean total upheaval; in fact, it could be fitted in virtually seamlessly to the conventional working of a good punched card installation.

The system described worked successfully from its inception in 1958/59, and was still going strong at Letchworth in the early 1970s. It was only brought to an end by a fire in the computer room that destroyed the 1202 that was still running it.

One of the offshoots of the system was the family of production control software packages that ICT and ICL subsequently developed, because the Letchworth system provided a focus for the group of production control consultants which grew up within the company.

Castlereagh never assumed the major role predicted for it. Top management decided that further investment in Northern Ireland was too risky, and they were probably right.

One final point: the Letchworth production control system more than paid for itself. The cost of the 1201 computer itself and all the development costs were more than offset by the resulting savings in efficiency, particularly in the greatly reduced cost of inventory, and in the reduced necessity for overtime and bought-out working. I spent nearly 30 years involved in the design and building of in-house computer applications for various parts of ICL, and I'm pretty sure that we very rarely achieved anything like that straightforward financial pay-back. But it was always fun trying.

Editor's note: This is an edited transcript of the talk given by the author to the Society at the Science Museum on 9 November 2006. Hamish Carmichael can be contacted at hamishc@globalnet.co.uk.

June – a Home-Grown Operating System

Andrew Colin

Nowadays, students on programming courses just key their code into a personal computer and get an instant reaction with the touch of a button. It was not always so. In the 1960s there were no microprocessors or personal computers. Universities were serviced by huge mainframes, and students had to punch their programs into cards, hand the pack into reception, and then wait 24 hours or more to get the results back. This article describes a custom-built system developed to speed up this process.

In 1965 I moved to Lancaster University from the Institute of Computer Science at London, where I had gained plenty of experience using the ICL Atlas. One of my first tasks at Lancaster was to mount a computing course for students doing Mathematics. We had an IBM 1620 – an underpowered slow machine even by the standards of the day. But with a good operator driving the various stages of the process, the compilation and execution of a small Fortran program took about 40 seconds, which was fast enough to give a small class of students several ‘turn-rounds’ a day.

The following year the 1620 was replaced by an ICL 1909, a state-of-the-art computer several hundred times faster (on paper) than the 1620. The machine was equipped with an ‘executive’ program (called ‘Exec’) to handle the peripheral devices and organise multi-programming, and the George 2 operating system to provide an interface to users and operators alike. My computing class had grown in size considerably, and I was confident that the machine would provide an excellent service to the students. They should have as many turn-rounds as they needed, without lengthy periods of waiting.

The day came that the machine passed its acceptance tests and I got my hands on it. I was shocked to discover that the minimal job, which had taken under a minute on the IBM 1620, now took *three and a half minutes* to run on the new machine, with the operator’s teletype running at full speed all the time. The reason soon became obvious: compiling, linking and running a program required the opening and closing of various files on the magnetic tape drives, and each such operation was reported in detail on the teletype at 10 characters per second.

My first instinct was to phone the Software department at ICL to ask if there was any way to suppress all this unwanted typing. The reply was that there was no way to prevent it; didn't I realise that the information was essential to the smooth operation of the system?

ICL had a policy of secrecy about many details of its machines. In particular, no documentation or listings of the executive code was ever to be released to clients. Nevertheless, with the help of the local engineer (who must by now be retired, otherwise I would not write this) I got a printed listing of the Exec, and decompiled the section that looked after the console printer.

This knowledge helped me gain access to the Exec team at ICL, and they told me of the "Trusted Program" feature, designed primarily for the use of operating system writers within ICL. A trusted program enjoyed certain privileges. In particular, it could intercept and suppress console messages, it could handle clock interrupts, and it could load and run other programs within its own address space, and handle their calls for peripheral transfers.

These facilities were well designed to support an operating system. The system we built at Lancaster was called "June" which was the name of my secretary at the time. Our design was heavily influenced by the Atlas operating system.

June was totally different in character to George 2. It made no attempt to satisfy computing needs in a completely general sense, but was designed to handle simple jobs as fast as possible, with minimum operator intervention.

Users submitted 'jobs' which could be either on cards or paper tape. The first few cards or lines consisted of a job description, similar to the one used on the Atlas, with details such as:

- The job title (used to label the printed output in huge characters)
- The user name
- The language used (Algol 60, Fortran, Cobol, or Assembler)
- Upper limits for the computation time and lines of output (this to prevent looping programs or wallpaper generators from clogging the system indefinitely)
- The identity of any magnetic tapes needed (not commonly used as most jobs did not need them).

Most of these details had standard defaults, so that the minimal job description had only a title and a name.

Each job was terminated with a card or line bearing three stars.

Once the system was running, the operators had only two tasks:

- Load each job (or each batch of jobs) into the card or tape reader as soon as it arrived
- Remove output from the line printer, burst it into separate jobs, marry them with the corresponding input media, and return them to the users.

The system was based on four concurrent processes

1. Card input: As long there were cards in the reader they were read and copied to a buffer on magnetic tape, called the 'input well'.

2. Paper tape input: Similarly for paper tape.

3. Job handler: This process used a scheduler to select a job from the input buffer, and run it by interpreting a simple script for each language. There was no attempt at (and no need for) multiprogramming; the system processed only one job at a time.

The system contained a few preset scripts, one for each language that the system supported. The script for a job specified a simple chain of programs: first a compiler for the appropriate language, then a linker (called a 'consolidator' in those days) and finally the object program itself. The job was aborted if either the compiler or the linker reported a failure such as a syntax error or unfilled reference.

As the job was running, the operating system intercepted all calls for input from cards or tape, and satisfied them from the data in the input well. Similarly, results intended for the printer were caught and sent to another buffer called the 'output well'. Any job that tried to print more than its print allowance would be stopped.

The operating system received regular clock interrupts, and stopped any job that ran over its time allowance.

4. Line Printer Output: This process handled the output well, buffering up a stream of data from users' programs and sending it to the line printer.

The four processes were synchronised using primitives available to trusted programs.

The Scheduler tried to give priority to short jobs; but longer jobs that had been in the queue for some time were gradually upgraded, so that they would eventually be executed.

At the time we were very conscious of the high cost of computing. At the end of each job the system printed an account of the total resources used – input, output and computation time, and an overall cost (in Pounds, Shillings and Pence!).

The system was designed and written over about six weeks. It was successful – it could handle small student jobs in about 12 seconds each. Students still had to hand their jobs in to the reception area, but got the results back in 15-20 minutes. The system was fully able to handle the larger jobs submitted by staff and postgraduates.

The total code size of the system was 2766 instructions, which left enough space in our 16K memory for object programs to run without too much memory swapping.

The system developed for some years.

When the machine was upgraded with an exchangeable disc, we transferred the input and output wells to the disc. The minimum job time fell to seven seconds.

We connected some teletypes to the system and built a primitive timesharing system. At that time, the cost of the necessary hardware was so high that very few people could take advantage of this feature.

I attended a conference on Operating Systems, where I met some of the George 2 designers. To my surprise, they were completely unaware of the excessive time their system took to process small jobs. It emerged that ICL's policy was to keep programmers and operators strictly apart. These people had never actually seen George 2 run, and simply had not realised that every tape operation was reported in such tedious and time-consuming detail.

The June system and its variants continued in use for several years, until it was made obsolete by the advent of personal computers for teaching programming. With hindsight, June was a good solution to a problem strictly of its own time.

Editor's note: Andrew Colin can be contacted at Andrew@crm.scotnet.co.uk .

Memories of the Elliott 803B

Graham Phillips

In 1961-62 I commissioned five Elliott 803B computers for Fairey Aviation, Brunel University, a Saudi Arabian bank, the BBC and a Welsh steelworks. I also commissioned paper tape stations (one or two Elliott paper tape readers, one or two Creed paper tape punches and a teleprinter) and an Elliott 35 mm magnetic film controller with four handlers. Two of the 803Bs were fitted with floating-point arithmetic boards and one was fitted with square root boards.

The computers employed ferrite core logic with transistor amplifiers and ferrite core storage (Mullard or Plessey 4K or 8K to order), and the power supply had a nickel-cadmium battery floating across the DC supply to provide smoothing and a temporary power supply in case of mains dropout. I believe the battery was of a type designed for use in aircraft 24V 400Hz systems.

The machine also included about eight hardwired initial orders, which meant that it was a self-starting machine. Another advantage of the initial orders along with using the B-line bit in the instruction code was that a sequence of instructions could be appended to a paper tape program so that the tape stopped in the reader at the end of a program instead of hurtling onto the floor. Paper tape was later superseded by plasticised (mylar) tape.

In those days the printed circuit boards were pre-tested only to the extent of checking and clearing the short circuits left between the DC power tracks by the flow-solder machine. The backplanes of the computers and peripherals were hand soldered and usually included many dry joints.

Commissioning meant taking cabinets of junk and getting them to work well enough to pass 24-hour soak tests at specified temperature limits. Because the computer had a loudspeaker attached to its instruction register it 'warbled' as it ran repetitive programs, and it was a thrill to come into work on the morning after a soak test and hear the machine warbling the welcome news that it hadn't failed in the night. I was also fascinated by the physics of the magnetostriction delay lines.

During the time that I worked at Elliott Brothers, Autocode was introduced with a fanfare as the way to make computers easy to program.

The 803Bs were previously programmed using an octal code (ranging from 00 to 77), which we commissioning engineers continued to use when we wrote test programs.

Minilog technology is sometimes confused with ferrite core packages, but in fact they were different things. Minilog packages, employed in the paper tape stations and the film controllers and handlers, consisted of resistor transistor logic and were much larger than the ferrite core packages.

Some people are unaware that there was an 803A machine whose sequence control register, instruction register and accumulator circulated in a continuous loop. I guess that it was too slow to be commercially viable.

The 803B was always spoken of as an Elliott Brothers computer, not an Elliott Automation computer. Elliott Automation produced Minilog and relay control equipment for industrial processes and the like. Its production floor was situated a short distance away from the 803B floor but in a different building.

Some of the terminology was different from that used now. For example, what are now usually called jump instructions were called transfer instructions, and what is now usually called the program counter was called the sequence control register. Most interesting was that a one-bit store, which had been called an Eccles Jordan circuit or a bistable multivibrator in valve technology and is now known universally as a flip-flop, was called a staticisor!

When I started on the production line we commissioning engineers were using Cossor oscilloscopes of a pre-war design. You can imagine how delighted we were when they were replaced with Tektronix oscilloscopes that enabled us to use a delayed sweep display to inspect every single pulse in a train of pulses (the 803B was a serial machine) and to measure time and amplitude accurately. I carried a 'button' battery in my pocket, which I used as a reference for calibrating the oscilloscope.

I remember well the enthusiasm of everyone involved. For example, I trained a service engineer and we used to work well into the evenings for no extra pay. To show his appreciation he took me for Indian meals in St. Albans, the first I'd ever had. The waiter was an Indian mathematics student who was studying for a degree.

Another enthusiast, a chap called Vic, was the chaser responsible for keeping the line supplied with components. One day a commissioning engineer needed a component which was out of stock so Vic immediately

jumped into his car and drove around the area till he found a shop that had one.

My single most pleasing moment occurred in Northampton. Three service engineers had tried without success over a period of several weeks to get an 803B that was installed in the Northamptonshire council offices to work. The local council official responsible for the machine had threatened to order Elliott to remove it from the premises if it wasn't available for use on Monday morning. On the preceding Friday afternoon I was asked to go and look at it. I grabbed my 35mm film can of test programs and because I didn't have a driving licence I was driven to the site by a mechanic whose usual job was servicing electromechanical peripherals such as punches and teleprinters.

When I got there I ran the soak test program tape, which passed through the reader and jumped out onto the floor. I ran successively simpler programs till I found one that failed but which was short. I discovered that in this simple multiplication program (101010... x 101010...) a conditional transfer instruction was transferring unconditionally. I put a board on a 'stalk', found a transistor with a base emitter short circuit, changed it and the machine was back 'on the air'. The mechanic looked at his watch and said, "We've only been here five minutes". He had been instructed to talk me into staying overnight and to book us into a hotel. Instead we had a slap up dinner at a Chinese restaurant and drove back to Borehamwood. Incidentally, it was the first time I ever travelled on the M1!

Editor's note: Graham Phillips can be contacted at gp2045@hotmail.com.

North West Group contact details

Chairman Tom Hinchliffe: Tel: 01663 765040.

Email: *tom.h@dsl.pipex.com*

Secretary William Gunn Tel: 01663 764997.

Email: *william.gunn@ntlworld.com*

Letters to the Editor

Dear Editor,

I am moved by Peter Hall's and Reay Atkinson's excellent pieces in *Resurrection* 35 to offer a salesman's-eye view of Government procurement in the mid-1960s. I was "Government Sales Manager" for English-Electric-Leo from late 1963 to 1967, which was mostly in the Wilson/Cousins era and before the "larger than Atlas" policy.

As Reay says, purchase decision-making for data processing applications at that period rested (*de facto* if not *de jure*) with project managers, ranking from SEO to SCEO. All, or nearly all, of them were highly competent, and most projects at that time succeeded.

A project manager's prime concern when deciding what equipment to buy was, obviously and rightly, that his project should succeed. He was also well aware that there was a preference for buying British. So far as I know this was never laid down as an instruction, but it was clear to project managers that they would be spared some hassle if they complied. The hassle was real: in at least one case the recommendation had to go to Cabinet more than once for a decision.

Most of them were in fact happy to comply. Both EE-Leo and ICT could provide suitable machines for most of the projects, and Treasury O&M and Technical Support Unit could and did give reassurance that the risk of going to either was low. A few determined project managers nonetheless forced through decisions in favour of IBM, Univac or Burroughs. This was mainly I think through loyalty to the supplier of a successful previous generation installation, rather than because the British machines were unable to cope, had inferior price/performance or fell short in quality of support.

In other words there was no explicit policy, but a more or less tacit tilting of the playing field. It actually worked rather well: there was still keen competition, the British industry was helped and the projects succeeded. I would be very interested to hear from former project managers whether this understanding of the situation is correct.

Incidentally, if Government projects succeeded then, why do they so often fail now? My crude answer would be because the 1960s projects were still essentially clerical replacements (albeit in some cases huge in volumes and complexity), without the ambitious system innovations (including making quasi-judgments) and large and often ill-disciplined

user populations of today. And is it possible that project managers are now *too* senior? Again I would be fascinated to hear informed opinions.

Yours sincerely,

John Aris

London W2

30 September 2005

Dear Editor,

I was interested in Hugh McGregor Ross's article in *Resurrection* 36, because I met him for a while when working for De Havilland Engine Co and later at Ferranti, Portland Place. I think I can add a little to the technical applications he mentions.

First, allied to aircraft design applications were those required in the design of jet engines. Following my attendance on what I believe was the first Pegasus programming course at Portland Place in 1955 (taken by Peter Hunt and George Felton) I wrote programs for engine design for which we used the Pegasus at De Havilland Aircraft Co, Hatfield. This was another example of using a computer to perform well-defined calculations in much shorter time than that hitherto been done by Brunsvigas.

Probably of greater interest though was a project in conjunction with Ferranti, Edinburgh, to produce tapes for driving numerically controlled milling machines. The application at De Havilland Engines was to produce compressor and turbine blades. Briefly, minimum data was input to Pegasus, which then produced paper tape defining the precise path of a cutting tool. This was sent to Edinburgh where it was converted to magnetic tape which drove the milling machine. This conversion was done by a special purpose device called a Digital Differential Analyser.

Bernard Wood

Bushey, Herts

by email from bernardwood75@hotmail.com

2 November 2005

Dear Editor,

I refer to the article “Finding the Necessity for Invention” in *Resurrection* 36. At the top of page 16 we find the words “Experimental Statistics”: may I elaborate on this theme, please?

I joined the British Iron & Steel Research Association (BISRA), in the Ironmaking Division at their Park Lane HQ in London, straight from school in January 1954. My boss at that time was the late JM Ridgion who was writing programs, in machine code and for the Manchester Mark IV, for statistical analysis of blast furnace performance. In time Mr Ridgion discovered a better way of performing these comparisons by “Materials and Heat Balances” (Schumacher).

Although eventually this work was moved to the North East Coast Laboratories of BISRA, and university graduates were also employed, I continued to be involved in this work until the early sixties, and so became involved in running the “Balance” program on the Ferranti Pegasus at BISRA Battersea. In 1964 I joined Ferranti Ltd, but at Bracknell and as a technical author!

I have summarised things as much as possible in this letter and, while matters in the iron and steel industry are now very much diminished in comparison with even 50 years ago, I thought I should bring the work of Maurice Ridgion to your mind.

AGS Davies

Thirsk, North Yorkshire

14 November 2005

Editorial contact details

Readers wishing to contact the Editor may do so by email to wk@nenticnap.fsnet.co.uk.

Forthcoming Events

Every Tuesday at 1200 and 1400 Demonstrations of the replica Small-Scale Experimental Machine at Manchester Museum of Science and Industry.

Every day Guided tours and exhibitions at Bletchley Park, price £10.00, or £8.00 for children and concessions. Exhibition of wartime code-breaking equipment and procedures, including the replica Bombe and replica Colossus, plus 60 minute tours of the wartime buildings. Go to www.bletchleypark.org.uk to check details of times and special events.

26 June 2007 Real Time Club 40th Anniversary Dinner at National Liberal Club, Whitehall Place, London SW1 (Web site www.realtimedclub.org).

12-14 July 2007 BCS 50th anniversary conference at Bletchley Park (12 July) and BCS offices in Southampton Street, London (13-14 July) – see www.bcsat50.org for booking details.

23 September 2007 Enigma Family Festival and Enigma Reunion at Bletchley Park – see www.bletchleypark.org.uk for full details.

18 September 2007 NWG seminar on “Making Money While the World Slept”. Speaker Gerald Everitt.

16 October 2007 NWG seminar on the English Electric Deuce. Speaker Mike Wetherfield.

29 November 2007 NWG seminar on artificial intelligence. Speaker Professor Max Brammer.

Details are subject to change. Members wishing to attend any meeting are advised to check the Society Web site or in the Events Diary columns of *Computing* and *Computer Weekly*, where accurate final details will be published nearer the time. London meetings take place in the Director’s Suite of the Science Museum, starting at 1430. North West Group meetings take place in the Conference room at the Manchester Museum of Science and Industry, starting usually at 1730; tea is served from 1700.

Queries about London meetings should be addressed to David Anderson at cdpa@btinternet.com, and about Manchester meetings to William Gunn at william.gunn@ntlworld.com.

Committee of the Society

Chairman **Dr Roger Johnson FBCS**, 9 Stanhope Way, Riverhead, Sevenoaks, Kent TN13 2DZ. Tel: 020 7631 6709. Email: r.johnson@bcs.org.uk

Vice-Chairman **Tony Sale Hon FBCS**, 15 Northampton Road, Bromham, Beds MK43 8QB. Tel: 01234 822788. Email: tsale@qufaro.demon.co.uk

Secretary and Chairman, DEC Working Party **Kevin Murrell**, 25 Comet Close, Ash Vale, Aldershot, Hants GU12 5SG. Tel: 01252 683503. Email: kevin@ps8.co.uk

Treasurer **Dan Hayton**, 31 The High Street, Farnborough Village, Orpington, Kent BR6 7BQ. Tel: 01689 852186. Email: Daniel@newcomen.demon.co.uk

Science Museum representative **Tilly Blyth**, Science Museum, Exhibition Road, London SW7 2DD. Tel: 020 7942 4211. Email: tilly.blyth@nmsi.ac.uk

National Archives representative **David Glover**, Digital Preservation Department, National Archives, Kew, Richmond, Surrey TW9 4DU. Tel: 020 8392 5330. Email: david.glover@nationalarchives.gov.uk

Bletchley Park volunteers representative **Pete Chilvers**, 3 North Way, Potterspurpy, Towcester, Northants NN12 7PW. Tel: 01908 543179. Email: petechilvers@beeb.net

Chairman, Elliott 803 Working Party **John Sinclair**, 9 Plummers Lane, Haynes, Bedford MK45 3PL. Tel: 01234 381 403. Email: john@eurocom-solutions.co.uk

Chairman, Elliott 401 Working Party **Arthur Rowles**, 74 Nottage Crescent, Braintree, Essex CM7 2TT. Tel: 01376 325990. Email: rowles01@globalnet.co.uk

Chairman, Pegasus Working Party **Len Hewitt MBCS**, 5 Birch Grove, Kingswood, Surrey KT20 6QU. Tel: 01737 832355. Email: leonard.hewitt@ntlworld.com

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

Email: bombe@jharper.demon.co.uk

Chairman, Software Conservation Working Party **Dr Dave Holdsworth CEng Hon FBCS**, University Computing Service, University of Leeds, Leeds LS2 9JT. Email: ecldh@leeds.ac.uk

Digital Archivist & Chairman, Our Computer Heritage Working Party **Professor Simon Lavington FBCS FIEE CEng**, Lemon Tree Cottage, High Street, Sproughton, Suffolk IP8 3AH. Tel: 01473 748478. Email: lavis@essex.ac.uk

Editor, Resurrection **Nicholas Enticknap**, 4 Thornton Court, Grand Drive, Raynes Park SW20 9HJ. Tel: 020 8540 5952. Email: wk@nenticknap.fsnet.co.uk

Archivist: **Hamish Carmichael FBCS**, 63 Collingwood Avenue, Tolworth, Surbiton, Surrey KT5 9PU. Tel: 020 8337 3176. Email: hamishc@globalnet.co.uk

Committee of the Society (contd)

Meetings Secretary and Web Site Editor: **Dr David Anderson**, The Wozzles, 1 Oatlands Road, Boorley Green, Hants SO32 2DE. Tel: 0239284 6668. Email: cdpa@btinternet.com

Chairman, North West Group **Tom Hinchliffe**, 44 Park Road, Disley, Cheshire SK12 2LX. Tel: 01663 765040. Email: tom.h@dsl.pipex.com.

Peter Barnes FBCS, 10 The Broadway, Gustard Wood, Herts AL4 8LN. Tel: 01438 832906. Email: panda.barnes@btinternet.com

Chris Burton CEng FIEE FBCS, Wern Ddu Fach, Llansilin, Oswestry, Shropshire SY10 9BN. Tel: 01691 791274. Email: cpb@envex.demon.co.uk

Dr Martin Campbell-Kelly, Department of Computer Science, University of Warwick, Coventry CV4 7AL. Tel: 01203 523196. Email: M.Campbell-Kelly@warwick.ac.uk

George Davis CEng Hon FBCS, 4 Digby Place, Croydon CR0 5QR. Tel: 020 8681 7784. Email: georgedavis@bcs.org.uk

Peter Holland, 10 Broad Walk, Orpington, Kent BR6 7RZ. Tel: 01689 891874. Email: peterholland@care4free.net

Eric Jukes, 153 Kenilworth Crescent, Enfield, Middlesex EN1 3RG. Tel: 020 8366 6162.

Ernest Morris FBCS, 16 Copperkins Lane, Amersham, Bucks HP6 5QF. Tel: 01494 727600. Email: Ernest.Morris@btinternet.com

Dr Doron Swade CEng FBCS, 54 Park Road, Kingston-upon-Thames, Surrey KT2 6AU. Tel: 020 8392 0072. Email: doron.swade@blueyonder.co.uk

Point of Contact

Readers who have general queries to put to the Society should address them to the Secretary: contact details are given on the page opposite.

Members who move house should notify Kevin Murrell of their new address to ensure that they continue to receive copies of *Resurrection*. Those who are also members of the BCS should note that the CCS membership is different from the BCS list and so needs to be maintained separately.

Resurrection is the bulletin of the Computer Conservation Society. Copies of the current issue are available from the Secretary for £5.00 each.

Editor -- Nicholas Enticknap

Typesetting -- Nicholas Enticknap

Cover design -- Tony Sale

Printed by the British Computer Society

© Computer Conservation Society