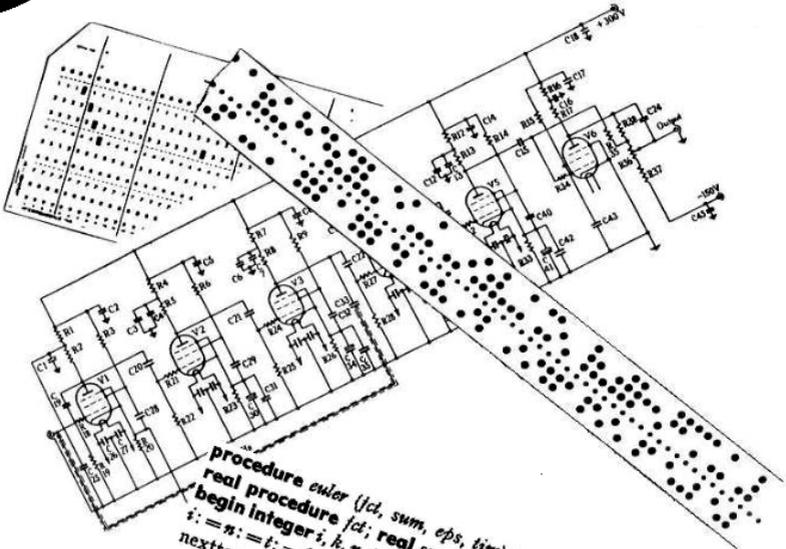


RESURRECTION

The Bulletin of the Computer Conservation Society



```

procedure euler (fct, sum, eps, tim); value eps, tim; integer tim;
real procedure fct; real sum, eps;
begin integer i, k, n, t; array m [0..15]; real mn, mp, ds;
i := n; t := 0; m[0] := fct(0); sum := m[0]/2;
nextterm: i := i + 1; mn := fct(i);
  for k := 0 step 1 until n do
    begin mp := (mn + m[k])/2; m[k] := mn; mn := mp end means;
    if (abs(mn) < abs(m[k])) then
      begin ds := m[n]; n := n + 1; m[n] := mn end accept
    else ds := mn;
  sum := sum + ds;
  if abs(ds) < eps then t := t + 1 else t := 0;
  if t < tim then go to nextterm
end euler

```





Computer Conservation Society

Aims and objectives

The Computer Conservation Society (CCS) is a co-operative venture between BCS, The Chartered Institute for IT; the Science Museum of London; and the Museum of Science and Industry (MOSI) in Manchester.

The CCS was constituted in September 1989 as a Specialist Group of the British Computer Society. It is thus covered by the Royal Charter and charitable status of 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,
- ◇ To develop awareness of the importance of historic computers,
- ◇ To develop expertise in the conservation and restoration of historic computers,
- ◇ To represent the interests of Computer Conservation Society members with other bodies,
- ◇ To promote the study of historic computers, their use and the history of the computer industry,
- ◇ To publish information of relevance to these objectives for the information of Computer Conservation Society members and the wider public.

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 BCS, fees from corporate membership, donations and by the free use of the facilities of our founding museums. Some charges may be made for publications and attendance at seminars and conferences.

There are a number of active projects 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 CCS also enjoys a close relationship with the National Museum of Computing.

Resurrection

The Bulletin of the Computer Conservation Society

ISSN 0958-7403

Number 66

Summer 2014

Contents

Society Activity	2
News Round-Up	12
Discovering an Enhanced Mercury <i>Gustavo Del Dago</i>	14
The HoneyPi Project – A Home-brew Computer with Attitude <i>Rob Sanders</i>	17
Strachey's General Purpose MacroGenerator <i>Andrew Herbert</i>	25
Obituary: Willam A. (Ben) Gunn (1926-2013) <i>Tom Hinchliffe</i>	31
40 Years Ago From the Pages of <i>Computer Weekly</i> <i>Brian Aldous</i>	33
Forthcoming Events	35

Society Activity

Planned Visit to Paderborn

We are planning a visit to Europe's largest computing museum, the Heinz Nixdorf Museum in Paderborn in Germany on 20th September. It will involve two nights in an hotel and longish journey (8 hours by train). If you are interested please contact Roger Johnson (r.johnson@bcs.org.uk) by July 31st.

Ferranti Pegasus — Rod Brown

As the Computer Conservation Society approaches 25 years of age I would like to report and reflect on the fact that one of the flagship projects from the Society's founding years was filmed for archive purposes by the Science Museum on the 5th June. The Ferranti Pegasus Project was started very early in the life of the Society and the filming event paid for by the Museum this year, is a reminder of just how successful it has been.



One of the early project leaders was John Cooper and if you care to locate the very first copy of *Resurrection* online you can read his report. However you should remember that this machine had been cared for by a succession of keepers, all of whom contributed to its survival many years before the CCS was founded in 1989.

So there was a lot to celebrate at the filming event. The day was a just a little tense for the filming crew and very tense for the team. As work had ceased on the machine some time ago we were not at all sure if enough of the machine would power up to allow filming to commence at all. Just to raise the tension a little more the Museum had closed the Mathematics Gallery for the day and isolated the smoke and fire alarms in the area.

The exercise was to record memories of its past users and to film the machine as it stands in the Museum today. My own abiding memory of those early days is of walking along Newman Street in London. It was the summer of 1960 and pausing by the window into the Ferranti showrooms I saw what I now think was

a Pegasus being used. How lucky I have been to be allowed to work on such a piece of history again at this end of my life!

The filming crew first undertook individual interviews with the team members. Eventually the moment could be postponed no longer and we were faced with trying to remember which power breakers needed to be turned on and just where they were located. But someone had to press the power up sequence button to bring the machine to a state in which it could be filmed.

Eventually the moment arrived and Chris Burton pressed the button with his right hand, holding aloft his left hand with fingers crossed. I am glad to report that after a few fuses were replaced, enough of Pegasus ran to allow the filming to be completed. Not all of the machine certainly, but enough.

The team (right) consisted of Glenys Gallagher, Peter Holland, Len Hewitt, Chris Burton and myself. Looking back at the succession of CCS members who have contributed time, effort and dedication to this project over the twenty five year period I knew I was in good company.



So well done Pegasus. The machine served us well on the day and we hope the final short film produced will not only celebrate the CCS commitment to working with the Science Museum but show that the CCS can and does run great projects in its area of expertise. This project alone restored a machine to working state and then delivered many years of demonstrations to the public in the Museum. Our thanks go to all for giving time and help on the day and crowning the success of the Project.

ICT 1301 — Rod Brown

The 1301 Project has now recovered all of the original punched card packs containing the complete software for the system. These are all of the engineering software tests, the standard subroutines, the assembler and early application software. The software is contained in over 50,000 80-column punched cards.

EDSAC Replica — *Andrew Herbert*

The principal development over the past quarter has been the set up of the EDSAC gallery at TNMoC, the erection of the 12 racks and the installation of the cabling for electrical power distribution.

Alex Passmore has completed the design of the power supply control unit and construction will begin shortly. A continued long-term concern is the provision of an adequate three-phase supply for EDSAC.

As can be seen from the photograph a considerable number of chassis are now in the racks awaiting power and interconnection.



To date we have designed, produced metalwork and issued components for the construction of 88 out of the total of 142 chassis required. 50

are known to be complete although not all are yet fully tested. In particular, 37 of the 41 storage regeneration units required are out for manufacture most of which are fully or partially completed. Chris Burton has overcome noise problems with early built versions and a modified design is being used for new manufacture and retrofitted to those already made.

We have a further 33 chassis designs to complete. Nigel Bennée is working his way through the arithmetic unit (or “the computer” as the pioneers called it). James Barr is working on main control and is about to start construction of the order decoding part. Andrew Brown and John Sanderson are looking at how the initial orders are transferred into store when the START button is pressed. Peter Lawrence is looking at clock and digit pulse distribution. Bill Purvis is exploring the I/O circuits.

Peter Linington has completed his research of suitable wire for building long delay lines and will shortly resume work on his experimental prototype with the new material.

Andrew Herbert is constructing the aerial gantry required to support signal cabling between racks. From photographs of the original this is made of paxolin strips supported on 2BA studding sticking up from the top of each rack.

There has been considerable debate about the specific EDSAC the Project wants to achieve as its first objective. 6th May 1949 is the date on which the machine ran its first program, but then there was further refinement as the machine settled into useful everyday service. Martin Campbell-Kelly has been busy researching the surviving operational memoranda in the Cambridge Computer Laboratory archives to understand the evolution of the machine. Based on this we have set March 1951 as our target: this will give us a machine compatible with the first edition of the famous Wilkes, Wheeler and Gill book on programming EDSAC. (We will still be able to run the early programs by using the later "initial orders 2" to load a copy of "initial orders 1" preceding a 1949 format paper tape.)

As the prospect of an operational machine draws nearer, the project management committee is beginning to plan how we will use the replica to convey the experience of being an EDSAC user and set the machine correctly in its historical context.

An exciting development in the past few weeks has been the discovery of a cache of EDSAC circuit diagrams. These were saved from a dustbin by John Loker when he joined the Cambridge Mathematical Laboratory as a young engineer to work on EDSAC 2. The drawings are of great use in confirming some of our designs for circuit elements and answering some open questions, although they have to be used with caution since they relate to EDSAC as it was in 1955/6 when much had been changed from the original.

Harwell Dekatron — *Delwyn Holroyd*

Last month we had problems with the power supply which ultimately turned out to be caused by a simple dry joint. In the course of fixing that problem we inverted the stabiliser unit in the rack to give access to the wiring, which promptly brought on two additional problems.

The first of these was found to be a never soldered connection on one of the valve anode caps, which must have been that way for 63 years! The second problem was a 10V sawtooth wave superimposed on the -55V supply, which we had seen during commissioning in 2012 just before the Reboot Event, but the problem had gone away when the control valve for the regulator was re-seated.

We then spotted and repaired a dry joint where two component leads were wrapped around a third. However to our surprise the sawtooth was still present and no longer intermittent! A closer inspection of the diagrams revealed one of the two components in question, a capacitor, had been added at Wolverhampton at some point. We are not quite sure what the intention was, but the actual effect of this capacitor was to inject noise from the 250V rail directly to the grid of the regulation valve for the -55V supply — thus amplifying it onto that supply. We removed the capacitor and the problem was solved. This was a rare example where a dry joint had turned out to be a good thing!

Our power supply investigations also revealed the two series regulator valves from which most of the positive rails are derived were not carrying an equal share of the load, and one was consequently heavily overloaded. The similar circuit for the negative rails incorporates small light bulbs on the anode caps which would reveal this issue, so we decided to modify the circuit and add bulbs to the positive side as well, and the faulty valve was replaced.

Apart from the power supply problems and some dust in the relays, operation has been generally reliable.

ICL 1900 Group — *Delwyn Holroyd*

Alan Thomson is progressing plans for a CCS event to mark the 50th anniversary later this year. He has been chasing down some contemporary ICT films and working with Rachel Boone at the Science Museum to identify any suitable 1900 artefacts that could be used for a display.

We were recently given a mystery paper tape which has turned out to be the fixed portion of a previously lost 1901A overlaid disc executive — E1DS. Now we just need to find the overlays!

Brian Spoor and Bill Gallagher's emulation group is in the process of rewriting some of their basic peripheral code. A copy of #XPLC (PLAN2 basic peripherals only compiler) has been recovered and, with the help of the consolidator #XPCC, has enabled code to be compiled and consolidated on their 1901A emulator under E1HS exec. They have also recovered #XMS1/2/3 City and Guilds teaching compilers/runtime. If anyone has any information about these compilers please get in touch.

An EDS 4/8 disc pack has been located in an archive at Plymouth University. TNMOC has an EDS 8 drive but no packs, so the pack is being exchanged for one of the museum's many EDS 80 packs. An investigation has been started into the feasibility of a partial restoration of the drive, with the goal of being able to spin the pack.

IBM Hursley Museum — Peter Short

We have been promised an 082 card sorter (right), which is stated to be in full working order albeit not yet at Hursley.

Last month we purchased a Selectric Composer on eBay for 99p, and collected it from Weymouth. Sadly the seller had binned the three boxes of type elements and the user manual before collection.

We have also just received another 1052 keyboardless golf ball printer as used as a console printer on S/360.

We have marked the 50th anniversary of System 360 with a small display of artefacts from this system. We have added suitable material to our website at hursley.slx-online.biz/mainframe.asp.



Following our appeal in *Resurrection 65*, Dik Leatherdale suggested we ask the BT Museum at Amberley if they have a spare telephone handset for our acoustic coupler. They kindly donated one, although it is physically a bit bigger than the coupler's pads. Dik then

suggested that our coupler might be a US version (and hence a different size). This proved to be the case.

In the US, a warehouse containing many items of IBM 360/370 kit has been discovered including a 360/40, the first 360 model to be completed. As this machine was designed at Hursley, we would love to acquire this rare machine. Photos at tinyurl.com/wh360.

Ferranti Argus 700 (née Bloodhound) Project — *Peter Harry*

Good progress has been made recently and it is hoped that the Bloodhound simulator will be fully serviceable in the next couple of months, assuming no major faults occur. We had a period of several weeks where new faults developed. That position has been reversed after we fixed a fault with the Argus 700GX processor's floating point unit so we are now back on track. The main challenge for us in resolving faults is that we have few spares and their condition is often not known. In this case we were fortunate that the fault was due to one inverter on a 74S04 TTL device which was easy to find. Ferranti produced a whole set of test programs for individual boards that can be loaded from tape. These would be very useful in establishing the status of individual boards but unfortunately we do not have any Argus 700 test programs. Does anyone?

The second focus of our recent fault finding activity was on the Argus 700's PeriBus. All peripheral and I/O activity runs through this bus which has two forms; Peribus for running in parallel on various I/O backplanes and SeriBus which is its serial form connecting the Argus 700 to any I/O chassis. Conversion between SeriBus and Peribus is carried out on each I/O unit of which there are three in the simulator. We were getting indicator lamps that were permanently 'on' when they should be 'off' which was found to be a faulty Serial to Parallel Converter (MP72). The MP72 was failing to read the first bit of an eight bit word, again due to a faulty 74S04 device. A number of faults resolved so far have been due to faulty inverter TTL devices.

Other activities included yet another Farnell power supply repair, a G24 5S, due to a bead Tantalum capacitor going resistive. In future all Farnell power supplies, as they are refurbished or repaired, will not only have all the RIFA filter capacitors (metallized paper) changed but also all the bead Tantalums. Our Weir SM300 power supply also failed. It was found to be responsible for causing mysterious reboots of the system as it supplies the Winchester disk controller board. The Weir power supply has been replaced and awaits repair.

We are always on the lookout for spares. Ferranti manufactured several variants of the Argus 700 but we use the Argus 700GX with two GL processors, circa mid 1980s build. Also; if there any CDC Wren I, 25Mb, Winchesters (ST506 interface) or Cipher Quarterback tape drives (QIC-02) around then we can give them a new lease of life.

Software — *David Holdsworth*

Leo III Intercode

The current state of play was presented at the Leo 60th Anniversary Reunion on 6th April. Now that the Intercode Translator is preserved in working order, attention has switched to the Leo III Master Routine which is written in Intercode, albeit in a somewhat unconventional way. We have produced a typescript of the Master Routine which we can process with our preserved Intercode Translator, and the output from this reports the same sizes etc. as are shown on the printer listing from which we started.

We now need the Master Routine Generator Programme (08004) of which we also have the listing. This is written in Intercode. Its copy typing is under way.

We are well on the way to preserving the cornerstone software of the Leo III in a working form.

Web Techology

In addition to presenting execution facilities on a web page we are post-processing source listings output by the Intercode Translator to provide hot links which enable following of the execution path. This is extremely similar to our facilities for KDF9 Usercode.

Brooker-Morris Compiler Compiler for Atlas

Since the last report on the Atlas emulators, most of our energy has been expended on getting Bill Purvis's emulator (Unix & Java) up to speed by comparing its behaviour with that of Dik Leatherdale's (Windows & C#). Along the way various useful additions were made to both. A trace file facility was added which is able to log each instruction obeyed together with its result to an agreed format. A checkpoint/restart facility was also devised (again using an agreed format) which enables comparisons to be made between the intermediate machine states obtained by the two emulators.

Dik reworked the DEFINE COMPILER subroutine of the Compiler Compiler so that it worked with the full version of the Atlas Supervisor and Bill then debugged a tiny "HELLO WORLD" program. That program has now been compiled, saved and run. Although it isn't actually a compiler, it has many of the characteristics of a compiler not least its relationship to the host operating system — the first time this has been achieved in more than 40 years.

Elliott 803/903 — Terry Froggatt

I visited TNMoC on 24th April with friends, and I was slightly embarrassed to find that the 903 was not working. Luckily, the fault was quickly located to the bit-slice register card in the +128 position, which I swapped with a card from the scrap 903. Some suspect transistors on the faulty card have since been replaced.

I was unsure as to whether the 903 extra store was working, so when I was at TNMoC again on 10th May, we ran the 16K version of BASIC. This ran for about 30 minutes before failing. It looks as though there is an occasional fault in the extra store which may prove hard to find.

The Elliott 803 was working normally.

Two members of the CCS committee visited the National Museum of Scotland's computer collection last year, and found two Elliott 903s and a matching desk there. I visited the collection on 1st April, specifically to look at these items in more detail. One 903 (from Heriot-Watt University) was in pieces and missing most of its cables although we did find some of the cables elsewhere in the collection. The other 903 (from the Royal Infirmary of Edinburgh) was in very good condition. It had clearly been moved intact with most, if not all, of its cables in situ. The third unit was identified as the electronics for a magnetic tape system, also from the Royal Infirmary, which appeared to be complete. The tape decks would be Ampex, but it's not clear if they were in the collection.

Regarding the 903 display unit cable, which the CCS has agreed to fund, I ordered the plug on 5-day delivery, only to be told that it would be delivered in late July. Then one arrived in early May but it was damaged. The UK supplier is currently trying to sort this out with the US manufacturer.

TNMoC — Kevin Murrell

TNMOC opening times for the general public have been extended. The whole Museum is now open to the public on Thursday, Saturday and Sunday afternoons from 12:00 to 17:00. There are also other opening times especially during school holidays publicised on the website. Colossus and Tunny continue to be open to the general public every day from 11:00 to 17:00. The Museum is open to pre-booked school and university groups and corporate groups daily.

In February a Cray 1 arrived from Farnborough Air Sciences Trust and is proving highly popular with visitors. Because it would require a Freon cooling plant that would occupy about twice the space of the existing machine restoration is impractical.

The disagreements over the fragmentation of Bletchley Park by the erection of gates and fences by the Bletchley Park Trust continue. TNMoC has restated its call for an independent review.

Differential Analyser — *Charles Lindsey*

Last time I reported on problems caused by a worn half-nut on our plotter (which we solved by cannibalising from an input table). We have now had a similar problem with the other axis of the plotter, and have obtained a (largish) quote for remanufacturing this component (mainly to demonstrate that remanufacture is feasible, and so allow us to continue regular demonstrations well into the future).

The machine is otherwise running well, and we are in process of bringing the Frontlah unit into regular use.

We are now turning our attention to installing the Inverter to allow variable speeds and reverse rotation of the main drive motor.

We are now promised access to the parts of the machine in offsite storage “within months”, a considerable improvement on “one of these years”, but still frustrating.

ICL 2966 — *Delwyn Holroyd*

The machine continues to run diagnostics successfully, and reliability is reasonably good.

A new volunteer at TNMOC who is an ex-ICL field engineer has started cleaning and inspection of one of the GTS2 magnetic tape decks with a view to restoration.

We have just received a donation of an ICL 7501 terminal, including the original furniture, floppy drive units, manuals, printer, spare boards and many discs. This will enable us to expand our existing 7501 configuration to include the 7700 word processing software.

We have now negotiated a museum licence with Fujitsu which will enable us to run VME on our recently donated Trimetra DY server. Work has started on getting the Trimetra DY running but unfortunately it appears to be suffering from disc problems.

News Round-Up

This coming November marks the 65th anniversary of CSIRAC, the pioneering computer built in Sydney, now believed to be the oldest extant stored-program computer in the world, albeit not working. Over the years many computers have been recorded as having been “retired” or “pensioned off” but CSIRAC will perhaps be the first reach official state pension age.

101010101

In his budget speech Chancellor George Osborne, announced the creation of *The Alan Turing Institute* which is to be a new research organisation to investigate methods of analysing huge quantities of “Big Data” to find patterns. A budget of £42,000,000 has been allocated to support this work. Organisations, including universities, will be invited to bid for the new institute later in the year.

At least two Alan Turing Institutes have been set up previously, one in Holland and another by the late Donald Michie in Glasgow in 1983 which latter closed in 1994 owing to lack of clients (a closure which seems to have escaped the notice of both *The Daily Telegraph* and of *The Guardian*).

Manchester MP John Leech immediately suggested that Manchester University would be an appropriate site for the new organisation, saying “Alan Turing’s contribution to Manchester was enormous...”. CCS members will, no doubt, have their own views about that.

101010101

In *Resurrection 65* we reported that Benedict Cumberbatch was to play the part of Alan Turing in the new film based on Andrew Hodge’s “magisterial” biography *Alan Turing: the Enigma*. More information has now reached us that the film is to be called *The Imitation Game* and is to be released on 14th November. Keira Knightly plays Turing’s one-time fiancée Joan Clarke, a casting which has attracted some criticism in that Knightly is rather too glamorous for the part.

It remains to be seen whether Cumberbatch’s researches into Turing’s voice has yielded benefits. Those in the know suggest that Derek Jacobi’s performance in Hugh Whitemore’s 1996 BBC dramatisation *Breaking the Code* was “spot on”. Judge for yourself, perhaps, at www.youtube.com/watch?v=EpSwP1XapTI.

101010101

The Imitation Game is, of course, another name for the *Turing Test*, Turing's proposal that Computers could be deemed capable of "thought" if a human in conversation with a machine could not distinguish the computer from a real human. On the 7th of May, exactly 60 years after Turing's untimely demise, it was announced that the Turing Test had at last been passed. In an annual event organised by the University of Reading at the Royal Society in London, *Eugene* a program simulating a conversation with a 13 year-old boy passed the stringent criteria set by the organisers. See tinyurl.com/ttsuc for more details.

Not everybody agrees however. One Chris Hughes wrote to *The Guardian* — *The Turing test has not been officially passed at all. Turing said that most of the interrogators had to be fooled, and that the conversation would have to take a long time. Plus, it's a chatbot, not an artificial intelligence program; and pretending to be a child whose first language is not English is clearly a cheat. AI is an impossible and wildly hubristic project. Give it up.*

101010101

The CCS FTP website at which a collection of emulators may be found has recently moved to www.cs.man.ac.uk/CCS/Archive/.

101010101

We regret to report the passing of Ian McNaught Davis, the genial host of BBC's *The Computer Programme* which gave life to the BBC Micro in 1982. Many warm tributes have been paid but tinyurl.com/imndobit is certainly the best. Ian attended our seminar on the BBC Micro in 2008 and proved to be a memorably entertaining participant.

101010101

CCS Website Information

The Society has its own website, which is located at www.computerconservationsociety.org. It contains news items, details of forthcoming events and also 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 www.cs.man.ac.uk/CCS/Archive/, where there is other material for downloading including simulators for historic machines. Please note that the latter URL is case sensitive.

Discovering an Enhanced Mercury

Gustavo Del Dago

In the course of a software preservation work, we have discovered that the Ferranti Mercury computer installed on the *Instituto de Calculo* at the *Universidad de Buenos Aires* had an improved memory management facility. We do not have the machine in order to verify this fact. What we do have are the oral memories of the protagonists and the recently discovered source code for a simulator for the IBM 1401 written in Mercury assembler language. This is a program with a high historical value for us. We know that that this cannot have been developed without the improved memory management facilities. Here we present a brief report of this work in progress with the intention of sharing our results. Any further information would be welcome.

Main Memory Organisation

The main memory or computing store had a capacity of 40960 bits. The memory can be addressed using short, half or long-words. It is important to note that each word size corresponds to a particular usage. The long-words (40 bits) are used to store floating point numbers, the instructions are located on half-words (20 bits) (called registers in the Ferranti literature) and the short-words (10 bits) are designed to store integers. Of course the programmer had some flexibility to store information in the main store but there are some restrictions too. The principal limitation is the size of the instruction's operand, which is 10 bits wide and so can only address locations in the range [0, 1023].

Given this addressing limitation the reader will appreciate that floating point numbers may occupy any part of main store but half-words may be located in the first half of the store and short words (in theory) the first quarter. However, the system offered some facilities in order to let the programmer use all of the first half of the memory for short-words. Each program instruction is assembled on a half-word and the Instruction Pointer Register is 10 bits wide. The program code must be accommodated in the first half of the memory space. In this scenario the program and the short numbers compete for the same memory space.

Ferranti's Solution

The "Input Routines" (IR) is a piece of software that implements some system level facilities. In the case of the Mercury the IR includes a machine code assembler. Some manuals refer to it as the PIG-2, but we can confirm the origin of this name. The assembler interprets different directives in order to read the

source program and generate the binary code. When the assembler found an instruction with a short operand it checked the value in order to determine if it corresponded to the first or second memory quarter. The machine implemented two different instructions sets, one for each of the two lower quarters of the store, and the assembler selected the correct code depending on the argument value. Figure 1 shows all the instructions with two codes. This mechanism offered a small abstraction to the programmer, who did not need to think about which quarter would be used to store each short number. The limitation therefore was a half of the memory space rather than a quarter. A curious aspect of the multiple codes used for some instructions, is that in two references: *the Instructions Chart* written in the IC (circa 1964) and in the *Programmer's Handbook for the Ferranti Mercury Computer* written by Ole-Johan Dahl & Jan V. Garwick, only one representation of the two instructions (first quarter memory code) is shown in binary. In the Ferranti Programmer's Handbook the codes are documented in decimal and octal. Here the functions with more than one code are marked by an asterisk. We suspect that it was not a big problem because the programmers had the IR in order to assemble their programs but the tables in other references are of help for entering instructions via the console keys.

The Instituto de Cálculo Improvement

Mercury was a scientific computer. The distinction between scientific and business oriented computers comes from a time when the computer's architecture defined its primary use. Scientific applications do a lot of floating point computation and here the Mercury was a powerful machine and one of the first to include hardware support for floating point operations (addition, subtraction and multiplication). Because of this predicted use, the machine had a memory addressing scheme that facilitated the manipulation of long-words, because a long-word location could accommodate a floating point number.

When the *Instituto de Cálculo* founders decided to acquire a computer, they knew that in order to solve mathematically-oriented problems a scientific computer represented the best choice. When the computer was installed in May 1961, the scientists at the *Instituto de Cálculo* used Mercury Autocode (a high level language) but soon a group of people started to use the Mercury assembler. This group, locally called *Grupo de Programación*, devised a new, widespread spectrum of opportunities in computer programming and realised the power of the machine in areas other than numerical computation. In fact this was the seed of the *Carrera de Computador Científico* (Computer Science School) created in 1963 at the *Universidad de Buenos Aires*. Two representative works are a Russian to Spanish translator and the COMIC language and compiler (a derivative of the Autocode). In this new scenario the need to process short

numbers became more frequent. A solution was offered by the *Computación Electrónica Group* (headed by Jonás Paiuk) who added two new machine instructions (functions 65 and 66). These instructions select the second half (function 65) or the first half (function 66) of the memory space. In effect this changed the base address of all the instructions which operated with short arguments. This new facility required explicit selection and it was the programmer's responsibility to know where is the base address pointing at any time. In this way the programmer could work with the complete memory space without restriction on the placement of short-words.

The *Instituto de Cálculo* machine was closed in early '70s and we have no documents about this hardware modification. We cannot be sure about the origin of this particular design. In the *Mercury Programmer's Handbook* functions 65 & 66 were assigned to Graphical Output operations *Open* and *Close Shutter*.

Figure 1

	Function	Code 1 st Quarter	Code 2 nd Quarter
0	$B' = Bt' = H$	0101001	1101001
1	$H' = B$	0100010	1100010
2	$B' = Bt' = B + H$	0101000	1101000
3	$B' = Bt' = B - H$	0101101	1101101
4	$B' = Bt' = B / 2 - H$	0101100	1101100
5	$B' = Bt' = B \& H$	0101011	1101011
6	$B' = Bt' = B H$	0101010	0101110
7	$Bt' = B - H$	1101010	1101110
20	$S' = St' = H$	0111001	1111001
21	$H' = S$	0110010	1110010
22	$S' = St' = S + H$	0111000	1111000
23	$S' = St' = S - H$	0111101	1111101
24	$S' = St' = S / 2 - H$	0111100	1111100
25	$S' = St' = S \& H$	0111011	1111011
26	$S' = St' = S H$	0111010	1111010
27	$St' = S - H$	0111110	1111110
60	$H' = t$	0000010	1000010
61	$H' = hs$	0010010	1010010

Gustavo del Dago has worked in software development for over 20 years. His passion for computing programming evolved into a lively interest in computer history and software preservation. He is a member of the *Proyecto SAMCA* (www.proyectosamca.com.ar) and can be contacted at gdeldago@gmail.com.

The HoneyPi Project – A Home-brew Computer with Attitude

Rob Sanders

This article is about the origin of and background to this project which is itself described on my embryo website www.honeypi.org.uk. At a recent CCS meeting someone suggested that it might be of interest to members, so here are the bald facts.

It may seem macabre to start a computer resurrection project with a death, but it is a fact that mother-in-law died and left us with her home to sell and therein lay the original problem. For several decades while I was working for a life assurance company I had acquired much of their cast-off electronic equipment to fuel my electronics hobby. Indeed amongst other things I could map the entire history of word processing with my collection of components from the earliest Flexowriter with its Selectadata dual speed paper tape reader through the Dataplex machine with its magnetic card drives, Honeywell Keytape machines with magnetic tapes and IBM Selectric typewriters, Wang word processor with hard and floppy disk drives and daisy-wheel printer, along with the Honeywell Level 6 and DPS6 minicomputers. Even now I can still lay my hands on policy clauses in Flexowriter code on paper tapes and corporate pension fund contracts on magnetic cards. But the problem was that much of my hoard was stored in a shed in mother-in-law's garden and it all had to be moved into our garage to sell the property. Now possibly I could pass a few working minicomputers on to a museum but over a thousand logic boards from Keytape machines were more of a problem. The HoneyPi Project was my solution.

The Keytape machines used Honeywell Series 200 components, so it occurred to me that I might be able to build a computer with them, maybe even one compatible with a Honeywell 200. I had started my IT career programming an H200 in the mid-'60s and still remember it with affection. Building the computer wouldn't solve my storage problem, not unless it was so like an H200 that I could give it away to someone as a replica when it was finished, maybe along with a large collection of spare parts. The peripheral devices might turn out to be a dual-speed paper tape reader and punch, two magnetic card drives and a daisy-wheel printer but that would just add to the appeal. Perhaps my idea of using magnetic tape drive vacuum/pressure components to enable it to drink a glass of water while whistling Land of Hope and Glory would have to go on the back burner though. No, HoneyPi wasn't started as a worthy reconstruction project but

as an exercise in recycling. Nevertheless it has developed since so I will attempt to promote it.

To my knowledge the last Series 200 machines in use, maybe in existence, were scrapped in the year 2000 by a company in Pennsylvania which had continued to use them until the millennium bug forced them to be decommissioned, not because they themselves were unreliable but just because the software was. Even the late Dr. William L. Gordon, chief designer of the H200, hadn't known that they were still in use then, as his daughter told me. The H200 was first marketed 50 years ago and knocked the then popular IBM 1401 off its pedestal, forcing IBM to stop dithering and launch their System 360 before they lost too many customers. So the H200 has its place in history. However, it deserves to be remembered for more than that in my opinion. It wasn't just a clone of the 1401 but a more sophisticated machine that just happened to be compatible with it in many respects. Dr. Gordon's daughter found in his personal papers a document written by him describing the thinking behind many aspects of its design and his own words put the idea of the H200 as the so-called "IBM 1401 Killer" into perspective. She plans to give the document to a museum but I won't pass on its contents without first getting her permission.

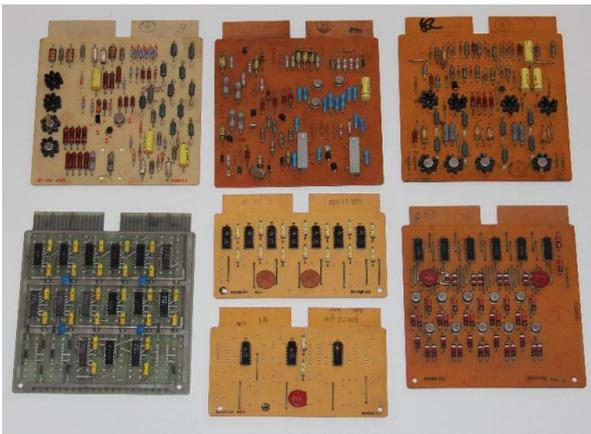
So, to assess the H200 in its own right, like the IBM 1401 it was a two address character stream processor in that one operation could process a string of characters of any length as



A Honeywell H200 – the brochure shot

defined by punctuation marks set in memory, but unlike the 1401 it used two bit punctuation marks which enabled it to recognise four level data structures made up of characters, words, items and records. For a newcomer to computers like myself in 1966 this was obviously what a computer ought to be able to do, but few machines have ever achieved it so simply at the hardware level in the way that the little second generation H200 did. It tackled data in such a naturally human way without the artifice of fixed length words that anyone could grasp the

principles and program it. It was a welcome change from working on spreadsheets with a hand-cranked calculator as I had been doing previously in the actuaries' department. A computer which treated all data just like numerical spreadsheets would hardly have caught my imagination in the same way, no matter how fast it was. Only later did I have to learn how to program machines that didn't give any consideration to humans. That said, its way of working wasn't far from the principle of a Turing machine with its data stream handling approach and, lacking in its simplest form a memory stack where data could be dumped thoughtlessly until needed later, it forced one to think about the machine's own state and use self-modifying code, machine state being another concept envisaged by Turing in his conceptual machine. Therefore despite the human-compatible nature of the H200 it taught me fundamental principles of computing which served me well in later years. I couldn't have found a better introduction to a successful career in software development, for the H200 was not a number-crunching computer nor a data processing machine but genuine information technology from the start. That was not surprising as, apart from being a modest general-purpose machine for smaller businesses it was also intended to be a front-end communications device for much larger people-unfriendly number-crunchers. Nowadays information technology focuses considerably on data stream handling and the H200 would be right at home. Finally, regardless of what went on inside the H200 it was a good looking machine when seen from the outside.



Various printed circuit boards

My heap of logic boards from seven Keytape machines mostly contained early integrated circuits from the 1960s with two gates in each package and (oddly I thought) connections for feedback capacitors. Possibly the feedback capacitors provided hysteresis which protected the gates against the high levels

of noise present on wire-wrapped backplanes. The H200 was a transistorised machine but I later discovered that at least one other computer in the 200

Series, the smaller H120, used ICs which was why it was also physically smaller. Therefore I am now justified in constructing my machine using those ICs. It is, after all, not a precise replica of an H200 but a pastiche, a machine in the style of the 200 Series, a machine that Honeywell could have built at the time. Apart from the logic boards I had a couple of backplanes with a total capacity of 200 logic boards between them, but my main problem was the magnetic core memory. Each Keytape machine had a single plane core memory capable of storing just a few punched card images. Furthermore, for the convenience of the operator they were wired for decimal addressing but the 200 Series used binary addressing. In his document Dr. Gordon explained the decision to deviate from the IBM 1401 precedent in this respect as he believed, wrongly as it turned out, that 1401 users only used addresses as meaningless field identifiers and didn't manipulate their values. Those who saw the H200 as a 1401 clone subsequently saw this aspect of its design as a weakness whereas users of the machine in its own right saw it as a benefit. In any case Honeywell had got the business approach right; to make a product similar enough to one's competitor's to gain customers but different enough to keep them. For my part I devised a way of mapping the binary addresses onto decimal ones in a non-sequential one-to-one way which could operate through simple logic in real time without slowing down the machine. I still had the problem of constructing a nine plane memory though.

Long ago I had given one memory board to a work colleague, as a memento, so I had left four planes containing 4200 bits and two containing 2100. It was impossible to modify the wiring of the memories without upsetting their noise cancellation properties so I would have to use them as they were. The only possibility would be to construct a 2k nine-plane memory by accessing the larger planes twice per byte, but that would slow the machine to well below the H200's operating speed and it would have a hard time outperforming the working 1401s still in existence. However, I still had one extra plane in hand, so I went back to basics and reviewed the way that core memories were used. Here I am considering the classic flux reversal mode of operation rather than the exotic differential technique as used in the H200 control memory. In his document Dr. Gordon explained that this control memory was the one extravagance in an otherwise frugal design and I believe that it was also the central reason for the H200 outperforming the 1401 as well as its for ultimate downfall. The control core memory ran at four times the speed of the main memory but that generated considerable heat and the differential storage method was heat sensitive, so if the ambient temperature was too high it would malfunction or even burn out, as my own employers found out the hard way. Honeywell claimed that the H200 didn't need a fully air-conditioned environment in which to operate and even demonstrated it in hotel rooms to prove their point, but as their company had

supplied environmental temperature control equipment from its earliest days I suspect that they made sure that the hotels they chose for their demonstrations



Another Installation

used them. No doubt as that was their key business they felt justified in assuming that their computers would always operate in such an environment.

The H200 was therefore very much a thoroughbred, reliable in familiar surroundings but unpredictable outside of

them. This may be the reason why none now exist, for any machine unwittingly operated outside of its comfort zone would eventually die. I do now have an original H200 control memory but have no intention of using it; there is a point where authenticity must give way to practicality. An engineer who used to work in the Computer Control Division at Honeywell (the continuation of the former company "3C" bought out by them) told me about the rivalry between the Series 200 engineers and his division. The Series 200 engineers regarded the CCD machines as "toys" while the CCD engineers regarded the 200 machines as unreliable dinosaurs which could easily be wiped out by just a little global warming. He gave training courses on the H716 to Series 200 engineers and to gain their respect he would turn off the H716 halfway through a demonstration and then turn it back on with the result that the machine continued working as though nothing had happened. Despite using core memory there was no guarantee what an H200 would do once it had lost power. The H716 also allegedly worked reliably at any temperature from freezing to 40°C. Nevertheless success with H200 sales justified Honeywell in calling themselves "The Other Computer Company" in their literature.

Returning to my problem, the classic method of using core memory distinguishes read cycles from write cycles and indeed multiplane memories with common write circuits for all planes and separate inhibit circuits for each plane virtually have to work this way, but my memory planes each had their own circuits, so I could use them how I liked. My solution was to use the fact that the sense amplifiers for a core memory always output the differences between what is being stored and what is already stored. Therefore one can effectively retrieve the existing contents while storing new ones by recombining the differences output with the inputs. That meant that I could process the memory in just three

cycles instead of four. Three cycles matched the speed of an H120, so my machine would still be compatible with a Series 200 machine. My technique involved retrieving the first half-byte in the first cycle, then storing it or its replacement in the second half-byte during the second cycle while simultaneously retrieving the second half-byte which would be stored or replaced in the first's location in the third cycle. This meant that bytes would continually switch between being stored big-endian and little-endian whenever they were accessed but I had a spare memory plane in which to store the current state of each byte. The key to the success of this approach was that it needed exactly what I had, two single-sized planes for the parity and state bits and four double-sized planes for the eight data bits. 50 years ago I might have applied for a patent but as it was I regarded it as a personal victory. I should also mention that 2k was the smallest memory size available in the original 200 Series, so my machine would still be within the specification. The project was feasible.



Core store module

I have mentioned this episode not for self-aggrandisement but because it illustrates a key aspect of such reconstruction projects. This was never meant to be a carbon copy of the construction of an H200 but a personal experience of resolving problems set in a similar context to that which pertained at Honeywell in the 1960s. One sees other examples of such reconstructions with people living in conditions as they were in the stone age or the 1950s and they are never entirely genuine. But they are undeniably educational and working out one's own solutions in such an artificial context is more rewarding than just slavishly following in someone else's footsteps. One revelation one can see in progressive versions of the Series 200 logic boards is the way that higher gate densities, brought about by integrated circuits, combined with the existing pin limitations of the standard board format forced the evolution of computer design as the specific structure of a particular machine moved out of the backplane wiring and onto customised printed circuits until ultimately the backplane was reduced to just a universal communications bus between complete sub-processors as in the level 6 minicomputers. For this project we must make similar decisions about whether we use many underpopulated general purpose logic boards or fewer high density boards made to our own specifications. Honeywell did it both ways, so we won't be outside our self-imposed guidelines whatever we decide. Anyway, now I could possibly build a machine to my satisfaction but there remained the

problem of whether one with only 2k bytes of memory could do anything worthwhile and also whether I could still program one. My answer was "The Pi Factory" as described in detail on my website, a program devised to meet constraints that no IBM 1401 could match. Not that I would think of challenging those worthies who keep the last of that breed working in America. Now with a plan and a demonstration program I did the unbelievable; I shelved the entire project for two years and wrote a novel.

My wife thought that I'd gone mad. I'd never even written a short story before in my life and she couldn't understand why I was in the spare room writing a novel instead of building a 1960s mainframe computer. Evidently when it comes to gauging insanity there is no norm for comparison. The truth was that, just like my HoneyPi project, it was just a pastime. Both the novel and the computer would just be by-products if either was ever completed. Writing the novel did tell me something though, first, that it was only the first story of a trilogy and perhaps life was too short to write the rest and build the computer. Second, that it wasn't just what one did but when one did it that was critical.

After a two year delay the errant muse that prompted me to start writing suddenly prompted me to return to the HoneyPi project with startling success. Through my meticulous searches on Google I made contact with Dr. Gordon's daughter and she was most interested in my plans. Her father had died a year or so earlier after suffering from Alzheimer's for five years and she'd realised since that she'd never paid enough attention to his boring perpetual enthusing over the H200. Corresponding with me prompted her to sort out his personal papers still stored in her basement and find that seminal document.

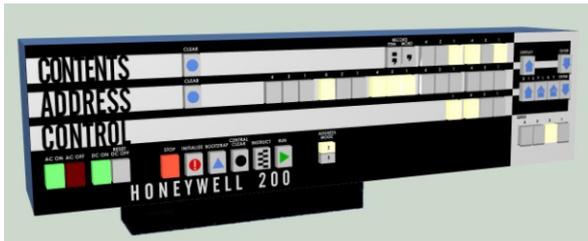
I also discovered Dr. Marcel van Herk, a relatively young vintage computer enthusiast who possessed at his home in Amsterdam almost all the components needed to build an original H200 memory unit that he'd acquired as a teenager. He was willing to entrust these to me for the project, so he became an active colleague with me in it and I, somewhat sadly, wouldn't have to implement my novel memory design after all.

Having all of Marcel's logic boards to accommodate meant that my backplanes wouldn't be adequate, so I had a new problem. As a contingency plan I had a foundry cast a replica of my existing backplane frame before we started using it but then another miracle occurred. Correspondence about EasyCoder, the H200 native assembly language, with a former programmer located in Australia resulted in someone else mentioning the engineer who maintained the last H200s in Pennsylvania in year 2000. It turned out that when he decommissioned them he'd taken home some 30,000 logic boards and around 30 backplanes to

sell to scrap dealers to help fund his daughter's forthcoming college fees. To cut a long story short Dr. Gordon's daughter in New Jersey was introduced to the last H200 site engineer in Pennsylvania and a backplane was shipped across the Atlantic to me while Marcel visited the said engineer's home during a trip to give lectures in the States and brought back a selection of extra logic boards.

I have made contact with a few former site engineers willing to search their memories for answers to pressing questions. I also found a manufacturer of printed circuit boards in Yorkshire who runs the business founded by his father round the year that the H200 was launched. So I now feel justified in having given my website that pretentious ".org" suffix.

At present I am busy building the 8k memory unit that Marcel and I agreed was the best option with our limited supply of boards. Even so Roger, our man in Yorkshire, will have to manufacture some replica boards at his factory to fill in the gaps once I determine whether Marcel's core memory modules are actually still functional.



The operating console

The final difficulty that we would have to face is that we don't have enough of those iconic chunky button tops to go on the control panel as illustrated on my website. Unfortunately a former work colleague

who operated our company's original H200 in the 1960s took its control panel home as a souvenir when the machine was decommissioned but threw it away when he subsequently moved house. Without that memorable control panel it just wouldn't be an H200, but the next miracle is apparently on its way. Someone has informed me that he owns an original H200 control panel, if he can find it. He believes that it is in the garden shed at his parents' home. Apparently the climate in California is much better for storing electronic equipment in garden sheds than in England and that must be an appropriate point to end this article.

In 1963 Rob Sanders passed over the opportunity to study pure mathematics at Cambridge and instead got career-long employment developing software for a City financial institution while Cambridge got Stephen Hawking to the lasting satisfaction of all concerned. Rob can be contacted at rob.s@honeypi.org.uk.

Strachey's General Purpose MacroGenerator

Andrew Herbert

As part of my on-going project to resurrect Elliott 903 software I regularly receive files from Terry Froggatt containing copies of paper tapes from his large collection. In one of the more recent batches there was a SIR assembly code tape labelled "GPM10 9/5/77, 900-Series Telecode"¹ and I guessed, correctly as it turned out, that this might be an implementation of Christopher Strachey's General Purpose Macrogenerator (GPM). This article gives an account of GPM; of errors found in the 1965 *Computer Journal* paper describing its use; a fatal flaw in the implementation given by the same paper; and an improved version of GPM developed fixing both problems.

GPM

GPM was a well-known and popular macro generator dating from the early 1960s GPM was initially conceived by Christopher Strachey when working as part of the Cambridge team designing the CPL programming language for the then new Cambridge Titan computer. Strachey proposed GPM as a tool to be used as an intermediate language for the compiler. The idea was to



translate CPL into GPM macro calls that would in turn be expanded into machine code. According to David Hartley this didn't work too well in practice — GPM was a more powerful tool than required for the task and was too easily abused, resulting in impenetrable code.

The GPM notation is deceptively simple. A string such as:

```
$DEF,Plus,<ADD ~1 ~2>;
```

¹ GPM10 can be found in the archive of Elliott software and documentation at: homepage.ntlworld.com/andrew.herbert1/andrew_herbert/elliott.html.

introduces a new macro "Plus" which when called with two arguments, **A** and **B** say, would generate:

```
ADD A B
```

Such a call would be written:

```
§Plus,A,B;
```

The symbols § and ; delimit a macro call. **DEF** is a built in macro that defines its second argument as having the value of its third argument. Commas separate the arguments of the call.

The symbol ~ followed by a digit n substitutes the nth argument of the current macro, if any. ~0 substitutes the macro name.

The symbols < and > are used to "quote" text – i.e., to stop it being expanded by the macroprocessor. The symbols are stripped off as the text is read, thus in the definition of **Plus**, the macro body is quoted to prevent argument substitution, but when **Plus** is called, the quotes are missing so argument substitution occurs.

The notation is completely general: macro definitions can be nested, arguments can include macro calls (including calls to **DEF**), quotes can be nested, etc, etc. There are a handful of built-in macros: **DEF** to define a new macro, **UPDATE** to replace a macro definition with a new body, **BIN** to convert a textual representation of an integer to binary form, **DEC** to convert a binary integer to an textual one and **BAR** which allows simple arithmetic on binary quantities. There is also a macro **§VAL, name**; that returns the unevaluated string bound to **name** as its result.

These simple facilities are sufficient to write quite complex programs. Here for example is a conditional:

```
§A,§DEF,A,<X>;§DEF,B,<Y>;;
```

This defines **A** to be **<X>**, then **B** to be **<Y>** and then immediately evaluates **§A**. If **A=B** the result is **Y**, otherwise the result is **X**. Thus we have a template for

```
if A=B then X else Y
```

From this we can progress to more exotic forms such as:

```
§DEF,Succ<0,1,2,3,4,5,6,7,8,9,10,§DEF,1,<~>~1;;>;
```

The effect of a call such as **§Succ,3**; is to make a temporary definition of a macro **1** and defining string **~3** and then call it with the arguments

2,3,4,...10, producing the result 4. So, the effect of `$$Succ,r;` where $0 < r < 10$ is to calculate the numerical successor of `r`.

There are further subtleties to do with how GPM handles superfluous arguments and the interplay between permanent and temporary macro definition for which the reader should refer to Strachey's paper.

The reader can perhaps start to see the veracity of David Hartley's observation that the GPM is perhaps too clever for its own good. That said, GPM was of considerable interest during its time as an object of academic study and is acknowledged as an important ancestor of the m4 macro generator used by Unix systems today.

Bugs

In his paper on GPM, Strachey develops a more complicated successor function

```
$$Successor,a,b,c;
```

that computes the c^{th} successor to the two digit number `ab`.

The definition is given as:

```
$DEF,Successor,<$~2,$DEF,~2,~1<,$Suc,>~2<;>;
```

```
$DEF,9,<$Suc,>~1<;,0>;>;
```

but when given to GPM10 the macro failed to work as expected and on pulling this loose thread it unravelled spectacularly. To work correctly the definition should have been:

```
$DEF,Successor,<$~2,$DEF,~2,~1<,>$Suc,>~2<;>;
```

```
$DEF,9,<$Suc,>~1<;<,>0>;>;
```

i.e., with some of the commas quoted, otherwise `DEF` is called with too many arguments and the wrong result produced.

This error was probably a simple transcription error from Strachey's original to the final typeset paper. Such errors are far less likely these days when we can cut and paste from program files to camera ready copy.

At first, when I found the GPM10 program wouldn't produce the expected result from the `Successor` macro, I suspected a bug in the code. Careful checking revealed the SIR code was a correct transcription of the CPL code given in the original paper, so then I began to doubt the CPL, especially as Strachey notes in the paper that the GPM implementation used at Cambridge was different to that specified in the paper. This led me to search the World Wide Web to see if there

were other implementations of GPM still extant. I discovered just one, in BCPL, maintained by Martin Richards, now a retired lecturer from the Cambridge University Computer Laboratory, a former colleague from my own time at Cambridge.

I learned that Richards had been Strachey's PhD student, hence his interest in GPM. Many readers will know Richards wrote BCPL as a simplified version of CPL while visiting MIT in the early '60s; that BCPL was a popular systems programming language in the '60s, '70s and '80s; and is an ancestor of C. Richards continues to maintain BCPL ² and in his collection of BCPL examples there was a program `bgpm` that claimed to be an implementation of GPM. Frustratingly `bgpm` also wouldn't run the **Successor** macro and worryingly GPM10 wouldn't run many of the test cases for `bgpm`. This led me to change tactics and code up my own **Successor** macro and in so doing I found the error in the original.

Looking at the `bgpm` it was apparent Richard's program was a different algorithm to Strachey's. In particular Strachey's CPL uses a single stack to evaluate macros and puts temporary macros on that stack using a linked list to maintain an environment chain that can be scanned to find the most recent definition of a given name. Richards by contrast uses two stacks – one for macro evaluation and a second for temporary definitions. I contacted Richards and he explained there was a fatal flaw in Strachey's algorithm: the special markers used to delineate environment chain information could be left on the evaluation stack after a call of **DEF** and give rise to errors (Strachey indirectly acknowledges this in section 2.6 of his paper). By using a separate stack for the environment, Richards sidestepped the problem. However in changing the data structures Richards had inadvertently changed the binding rules for temporary macros and this is what led to his examples not running on GPM10 and my proposed correction to Strachey's **Successor** macro not running on `bgpm`.

Following our discussion Richards wrote a BCPL equivalent of the CPL version in Strachey's paper. With this program, which included a useful backtrace facility and better error reporting than GPM10, I was able to confirm the error in **Successor** and its correction.

² `bgpm` & `bgpmcj` can be downloaded from www.cl.cam.ac.uk/~mr10/BCPL.html

bgpm Evolution

Following on from writing the BCPL version of Strachey's algorithm Richards went back and updated bgpm to handle temporary macros identically to Strachey, but importantly still with two stacks so that problems with environment chain markers couldn't arise. This version thus correctly expands all of Strachey's examples and, with some minor recoding, all of Richards' examples (including those which would have crashed Strachey's original due the environment chain marker problem).

Richards' bgpm contains several other improvements. In particular arithmetic is handled by a macro `eval` that understands simple numerical expressions. Thus

```
$DEF,x,1; ... $UPDATE,X,$BAR,+, $VAL,x;,$BIN,1;;;
```

becomes:

```
[def\x\1] ... [set\x\[eval\[x]+1]]
```

(Richards uses different delimiter symbols better suited to the ASCII character set.)

Richards' form is easier to read and importantly avoids the risk of leaving arbitrary binary values on the evaluation stack which might confuse the scanner (as does indeed occur with Strachey's algorithm if certain negative numbers are computed).

Two new macros `lquote` and `rquote` are provided that allow left and right quotes to be generated in the output stream, a completeness feature curiously missing from Strachey's original.

Argument naming is also tightened up. A `~` can be followed by any integer allowing arbitrary numbers of arguments, whereas in Strachey's GPM argument numbers greater than 9 had to be encoded using the symbols that followed digit 9 in the Titan character code, which is obscure.

Perhaps most importantly of all, bgpm allows comments to be inserted in GPM input. A ``` (grave) symbol causes the scanner to skip all remaining text on the current line and continue to skip white space until some other symbol is found. This is valuable in allowing complex macros to be documented inline and layout be used to improve readability without impacting adversely on the layout of the output.

Elliott 903 GPM

Having resolved the problems running Strachey's examples and working with Richards on the convergence of bgpm and Strachey's semantics, I wrote a translation of bgpm in SIR as a replacement for GPM10 and this is available from my Elliott 900 Series software and documents archive to any who might be interested. (Writing a 903 code generator for BCPL is a project for another day.)

The origins of GPM10 itself are unclear: it was not issued by Elliotts — there were GPM-like facilities in the Elliott MASIR (macro SIR) assembler, but not the complete set. And indeed why GPM10? I recall there being a GPM program for the 903 at Leeds University when I was a student there from 1972-5. The GPM10 tape however is dated 1977. I could easily imagine there were several versions in circulation at various times given the popularity of GPM in those days and the suitability of implementing GPM as a final year student project. If any CCS member can shed further light on these matters I'd be pleased to hear from them.

Summary

An early and influential macro generator has been preserved ³ with two implementations, one in Elliott SIR, the other in BCPL which, although now a vintage programming language, is still supported on most current systems including, most recently, the Raspberry Pi. During the preservation a number of bugs in the implementation and illustrative examples have been corrected, offering a much-improved version for the modern user. *The Computer Journal* has published a record of the errors on its website.

Following an exciting and senior computing career Andrew Herbert is now Project Manager for the EDSAC Replica Project at TNMoC. He can be contacted at andrew@herbertfamily.org.uk

North West Group contact details

Chairman	Tom Hinchliffe:	Tel: 01663 765040.
	Email: tah25@btinternet.com	
Secretary	Gordon Adshead	Tel: 01625 549770.
	Email: gordon@adshead.com	

³ Taking David Holdsworth's definition of "software preservation," namely being able to run the software on current platforms.

Obituary: Willam A. (Ben) Gunn (1926-2013) *Tom Hinchliffe⁴*

Ben Gunn led life to the full and put a vast amount of energy into everything he did. He died peacefully on 13th September 2013. He was the inaugural secretary of the CCS (NW), an office he held for 17 years until he and his wife Ivy moved to Yorkshire to be closer to their daughter Hilary. He was awarded a BCS certificate and an engraved decanter for his services to the CCS (NW), in 2011.



Ben began work at 16 when he left Redhill Technical College and took a job with the Post Office laying telephone cables in London during the bombing.

Ben never stopped learning and for the next 10 years, interrupted only by his time in the Army, he attended evening classes at Borough Poly and Regent Street Poly in London. He finished his time there teaching HND electrical engineering at Regent Street.

Ben started his IT career in the fledgling British computer industry in the late 1950s with Powers Samas in Whyteleafe, Surrey and remained there when ICT was formed. He then moved into the new R&D buildings in Stevenage, where he led teams working on the 1900 Series of machines, in particular peripherals. His

⁴ With help from Neil Gunn, Colin Skelton, Eva Hinchliffe and Gordon Adshead

passion was the early OCR (optical character recognition) machines used for cheque reading and led a team developing a reader for the then new OCR B font- much more readable than the old OCR A still found on cheques today.

His love of mountains was already established and he would frequently leave work on a Friday evening, drive to Snowdonia, camp, spend the weekend rock climbing and be back at work early Monday morning.

Small Mainframe R&D at Stevenage finished in the mid '70s and by that time Ben was in his late 40s when he moved again, this time to Manchester and West Gorton. At West Gorton he became a senior member of the teams developing the 2960, 2966 and DM1 mainframe systems.

The Welsh, Scottish and Lake District mountains were now much nearer his home which pleased him immensely. He was famously photographed reading a copy of *Computing* on top of the Matterhorn.

He was very involved with the community in High Lane where he lived and played an active role in most High Lane organisations. For example:- Abbeyfield, High Lane, sheltered accommodation for the elderly. Ben held various voluntary positions over the years, including Chairman. He was heavily involved in the initial planning and construction. The home was opened in 1984. And Ben worked strenuously over the years to make it a real success.

He was Chairman of the Mens' Forum for many years and an active member of the Residents' Association, working to get the best for the community. He was very keen on maths and for some years taught it at the U3A in High Lane.

Ben was a senior member of the High Lane Village Hall committee. After the original wooden hall was burnt down, he was involved in the planning and fundraising to build a new brick hall. This was fully achieved, and under Ben's leadership, events and shows for all ages were put on and it became very successful.

He was a keen climber and walker at home and abroad. He formed a friendship with a Sherpa on a trip to Everest and supported some of their projects in Tibet.

He enjoyed wood carving and made some good tactile carvings and complex mathematical shapes.

Ben was a well-liked man and always made a very positive contribution on everything he tackled. He will be missed by very many people.

40 Years Ago

From the Pages of *Computer Weekly*

Brian Aldous – TNMoC Archivist

FedNet fatal for privacy is how one US Senator sees the US General Administration's plan for a Federal Computer Network, claiming it will be the beginning of a massive data bank combining confidential information on all Americans into a single system. (CW 396 p24)

NPL and IRIA networks in Anglo-French link; The first test of 'host-to-host' protocols between the NPL's computer network and the French Cyclades network are expected to be made next week. (CW 397 p1)

Microprocessor bid by Ferranti, to stake a claim in the world microprocessor market, estimated to be worth £500 million by 1980, Ferranti is developing its own 'computer-on-a-chip'. Known as the F100L, it is expected to become available in 12 to 18 months. (CW398 p1)

BCL crash; Following the collapse of Business Computers Ltd. last week, users of BCL equipment have been urged not to panic. The receiver put in by the company's bankers is anxious users should not rush off to make their service and support arrangements. (CW398 p1)

Elliott 803B retires to college; Science and mathematics students at Braintree College to benefit from Elliott 803B being retired by Security Computing Services Ltd. (CW 398 p3)

London link to Arpanet enables otherwise impossible UK research projects to get off the ground. From the giant Illiac IV to the Data General Nova Mini, the network links over 60 host computers in the US and an IBM 360/195 in the UK. (CW 399, Int. Ed. 44 p4/5)

CWT rescue plans to save BCL business; In an eleventh-hour bid to allow the stricken Business Computers Ltd. to carry on trading, the receiver has reached agreement with Computer World Trade under which the field maintenance and computer broking group will continue all aspects of business essential to existing BCL users. (CW 400 p1)

DHSS plans countrywide network; Several million pounds' worth of equipment orders are expected from DHSS shortly as the programme for computerising the assessment and payment of social security benefits gets under way. (CW 400 p48)

ICL on-line system for NZ PO; The New Zealand Government has confirmed that a \$NZ3m contract to establish a nationwide on-line system on behalf of the New Zealand Post Office Savings Bank has been awarded to ICL. (CW 401 p40)

NCR launch 725 mini in UK; Widely used at NCR PoS installations in the US, and ordered in the UK by the John Lewis Group, the NCR 725 in-store minicomputer for controlling NCR 280 PoS terminal arrays has now been officially launched in the UK. (CW 402 p17)

GEC-Elliott traffic contracts; Presenting a breakthrough in a highly competitive market, GEC-Elliott Traffic Automation has won two contracts, worth £600,000 in all, to supply GEC 4080 computer-based traffic control systems to Wolverhampton and Northampton. (CW 403 p1)

Honeywell Series 60 £1m orders; A plethora of UK orders for its new Series 60 computers was announced this week by Honeywell. The £1m mark has now been passed with £988,000 of equipment being added to the £50,000 order for a 10K 61/58 from ASWE. (CW405 p32)

PDP-8/E gets mass memory storage; A mass memory system for the DEC PDP-8/E has been announced by Vermont Research of Leatherhead, Surrey. Known as the 3402 system, it follows the 3401 drum memory for the PDP-11 range announced earlier. (CW 406 p7)

Fairchild launch FIFO memory; A first-in, first-out memory from Fairchild Camera and Instrument Corporation has been designed to provide a low-cost solution to the problems of interfacing digital systems with different data rates. (CW 406 p11)

Argus minis in £1m Met Office order; Contracts worth over £1m have been placed by the MoD with the military systems division of Ferranti to supply 17 sets of ground station equipment and one calibration plant to the Met Office. Each unit ordered includes an Argus 700E minicomputer. (CW 406 p40)

DEC launch kit version of PDP-8/E mini; A kit version of the well-known PDP-8 minicomputer has been introduced by DEC. Called the PDP-8/A kit, it sells for £448 or less in quantity, a batch of 50 selling for £296 each. (CW 407 p29)

First 2903 bureau in operation; The first ICL 2903 bureau has been set up by a first-time user of an in-house computer. The user is part of the System Aid Group, Teltours of Southall, which handles travel accounting and documentation for tour operators. (CW 408 p28)

Forthcoming Events

London Seminar Programme

- | | | |
|---------------------------|---|----------------------------|
| 11 th Sep 2014 | The Origins of the ICT 1900 Range (1961-64). | Alan Thomson & John Buckle |
| 26 th Sep 2014 | Return visit to the IBM Museum at Hursley Park
100 years of IBM
The Years at Hursley Park
Only available by prior booking at https://events.bcs.org/book/1105/ . A modest charge of £12 for lunch applies. | Terry Muldoon
David Key |
| 16 th Oct 2014 | Annual General Meeting (CCS 25 th Anniversary)
Computer Conservation and Museums: Fight or Flight? | Doron Swade |
| 20 th Nov 2014 | Rescue, Retention and Relevance: The Three R's of Software Preservation | David Holdsworth |
| 18 th Dec 2014 | Computer Films | Panel of Speakers |
- London meetings normally take place in the Fellows' Library of the Science Museum, starting at 14:30. The entrance is in Exhibition Road, next to the exit from the tunnel from South Kensington Station, on the left as you come up the steps. For queries about London meetings please contact Roger Johnson at r.johnson@bcs.org.uk, or by post to Roger at 9 Chipstead Park Close, Sevenoaks, TN13 2SJ.

Manchester Seminar Programme

- | | | |
|---------------------------|---|------------------|
| 16 th Sep 2014 | The EDSAC Rebuild Project | Andrew Herbert |
| 21 st Oct 2014 | Rescue, Retention and Relevance: The Three R's of Software Preservation | David Holdsworth |
| 18 th Nov 2014 | History & Development of the Internet | Nigel Linge |

North West Group meetings take place in the Conference Centre at MOSI — the Museum of Science and Industry in Manchester — usually starting at 17:30; tea is served from 17:00. For queries about Manchester meetings please contact Gordon Adshead at gordon@adshead.com.

Details are subject to change. Members wishing to attend any meeting are advised to check the events page on the Society website at www.computerconservationsociety.org/lecture.htm. Details are also published at in the events calendar at www.bcs.org.

Museums

MOSI : Demonstrations of the replica Small-Scale Experimental Machine at the Museum of Science and Industry in Manchester are run every Tuesday, Wednesday and Sunday between 12:00 and 14:00. Admission is free. See www.mosi.org.uk for more details.

Bletchley Park : daily. Exhibition of wartime code-breaking equipment and procedures, including the replica Bombe, plus tours of the wartime buildings. Go to www.bletchleypark.org.uk to check details of times, admission charges and special events.

The National Museum of Computing : Thursday, Saturday and Sunday from 13:00. Situated within Bletchley Park, the Museum covers the development of computing from the wartime Tunny machine and replica Colossus computer to the present day and from ICL mainframes to hand-held computers. Note that there is a separate admission charge to TNMoC which is either standalone or can be combined with the charge for Bletchley Park. See www.tnmoc.org for more details.

Science Museum :

There is an excellent display of computing and mathematics machines on the second floor. Other galleries include displays of ICT card-sorters and Cray supercomputers. Highlights include Pilot ACE, arguably the world's third oldest surviving computer. Admission is free. See www.sciencemuseum.org.uk for more details.

Other Museums : At www.computerconservationsociety.org/museums.htm can be found brief descriptions of various UK computing museums which may be of interest to members.

Contact details

Readers wishing to contact the Editor may do so by email to dik@leatherdale.net, or by post to 124 Stanley Road, Teddington, TW11 8TX. Queries about all other CCS matters should be addressed to the Secretary, Kevin Murrell, at kevin.murrell@tnmoc.org, or by post to 25 Comet Close, Ash Vale, Aldershot, Hants GU12 5SG.

Committee of the Society

Chair: **Rachel Burnett FBCS:** rb@burnett.uk.net
Secretary: **Kevin Murrell MBCS:** kevin.murrell@tnmoc.org
Treasurer: **Dan Hayton MBCS:** daniel@newcomen.demon.co.uk
Chair, North West Group: **Prof. Tom Hinchliffe FBCS FIEE CEng:** tah25@btinternet.com
Secretary, North West Group: **Gordon Adshead MBCS:** gordon@adshead.com
Resurrection Editor: **Dik Leatherdale MBCS:** dik@leatherdale.net
Website Editor: **Dik Leatherdale MBCS:** dik@leatherdale.net
Meetings Secretary: **Dr Roger Johnson FBCS:** r.johnson@bcs.org.uk
Digital Archivist: **Prof. Simon Lavington FBCS FIEE CEng:** lavis@essex.ac.uk

Museum Representatives

Science Museum: **Dr Tilly Blyth:** tilly.blyth@nmsi.ac.uk
Bletchley Park Trust: **Kelsey Griffin:** kgriffin@bletchleypark.org.uk
TNMoC: **Dr David Hartley FBCS CEng:** david.hartley@clare.cam.ac.uk
MOSI: **Sarah Baines:** s.baines@mosi.org.uk

Project Leaders

SSEM: **Chris Burton CEng FIEE FBCS:** cpb@envex.demon.co.uk
Bombe: **John Harper Hon FBCS CEng MIEE:** bombeebm@gmail.com
Elliott: **Terry Froggatt CEng MBCS:** ccs2@tjf.org.uk
Software Conservation: **Dr Dave Holdsworth CEng Hon FBCS:** ecdh@leeds.ac.uk
Elliott 401 & ICT 1301: **Rod Brown:** sayhi-torod@shedlandz.co.uk
Harwell Dekatron Computer: **Delwyn Holroyd:** delwyn@dsl.pipex.com
Computer Heritage: **Prof. Simon Lavington FBCS FIEE CEng:** lavis@essex.ac.uk
DEC: **Kevin Murrell MBCS:** kevin.murrell@tnmoc.org
Differential Analyser: **Dr Charles Lindsey FBCS:** chl@clerew.man.ac.uk
ICL 2966/1900: **Delwyn Holroyd:** delwyn@dsl.pipex.com
Analytical Engine: **Dr Doron Swade MBE FBCS:** doron.swade@blueyonder.co.uk
EDSAC: **Dr Andrew Herbert OBE FREng FBCS:** andrew@herbertfamily.org.uk
Bloodhound Missile/Argus: **Peter Harry:** pdh@imttx.co.uk
IBM Hursley Museum: **Peter Short:** peters@slx-online.biz
Tony Sale Award: **Peta Walmisley:** peta@pwcepis.demon.co.uk

Others

Prof. Martin Campbell-Kelly FBCS: m.campbell-kelly@warwick.ac.uk
Peter Holland MBCS: p.holland@talktalk.net
Pete Chilvers: pete@pchilvers.plus.com

Point of Contact

Readers who have general queries to put to the Society should address them to the Secretary (see page 36 for contact details). 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 BCS, however, need only notify their change of address to BCS, separate notification to the CCS being unnecessary.

Resurrection is the bulletin of the Computer Conservation Society.

Editor — Dik Leatherdale

Printed by — BCS the Chartered Institute for IT

© Computer Conservation Society