

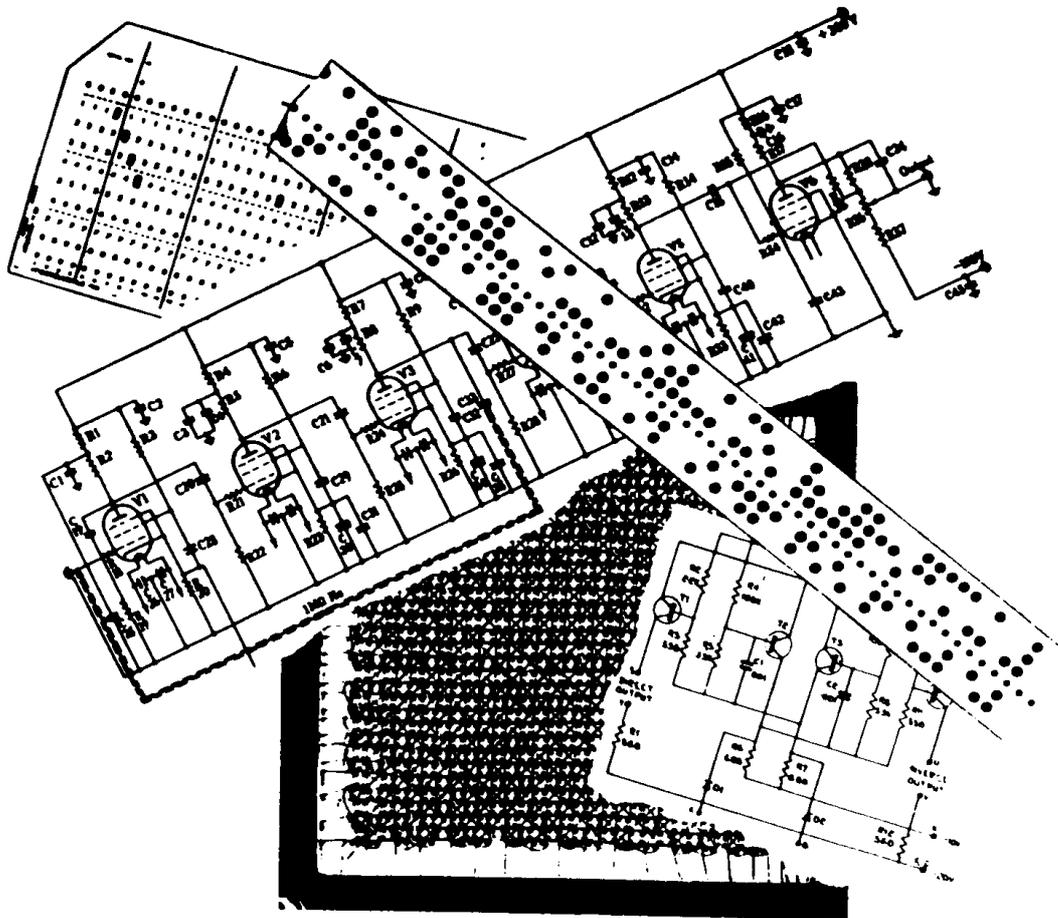
Issue Number 33

Spring 2004

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 33**

**Spring 2004**

## Contents

<b>News Round-Up</b>	<b>2</b>
<b>Society Activity</b>	<b>3</b>
<b>Inmos and the Transputer (part 2)</b> <i>Iann Barron</i>	<b>5</b>
<b>A Management Perspective on Orion</b> <i>Peter Hall</i>	<b>15</b>
<b>Before Moore's Law</b> <i>Brian Wichmann</i>	<b>19</b>
<b>Obituary: John Gosden</b> <i>David Caminer</i>	<b>25</b>
<b>Letters to the Editor</b>	<b>26</b>
<b>Forthcoming Events</b>	<b>27</b>
<b>CCS Web Site Information</b>	<b>27</b>

---

## News Round-Up

---

Rod Smith, who has been acting as liaison officer between the Society and the Science Museum for the past 10 years, retired in January. He was warmly thanked by the Committee for his untiring work on behalf of the CCS.

101010101

We look forward to work proceeding again on the Elliott 401 now that Arthur Rowles has become the new chairman of the Working Party in succession to Chris Burton. Arthur is a former employee of the Science Museum, who has the distinction of having recruited Doron Swade to the Museum back in 1975.

101010101

Arthur invites all members who have previously offered to join the 401 Working Party, or who would like to join it now, to get in touch with him. Contact details are inside the back cover.

101010101

Tony Sale reports that a Colossus cracked a code just before Christmas for the first time since 1945! Tony and his team are beavering away in a bid to ready the machine for the celebrations on 1 June marking the 60<sup>th</sup> anniversary of the first running of a Colossus Mark II.

101010101

The National Computing Collections Listing Project Web site is now accessible. Jenny Wetton reports that “it is fairly limited in the information contained within it and the quantity of collections covered, but members may want to consult it to find out what is held in the main public collections in the UK”. Interested readers can access it at <[www.sciencemuseum.org.uk/ncclp/welcome.htm](http://www.sciencemuseum.org.uk/ncclp/welcome.htm)>.

101010101

This year sees the Golden Jubilee of the Fortran programming language, which was first released in November 1954.

---

## **Society Activity**

---

### **Elliott 401 Working Party**

*Arthur Rowles*

Owing to commitment priorities falling elsewhere, little or no on-site work has been undertaken on the 401 at Blythe House over several years. Former chairman Chris Burton requested an alternative chairman be appointed; I am he. With only three days work under our belts to date it is much too soon to report progress; but a proposed programme of up to four on-site days per month appears to have been accepted by the small band of supporters, and additional manpower is being sought.

Specification details of a 401 Order Code Emulator System are being worked out by Alan Martin, with engineering under Microsoft Windows proposed. This work is directed at (a) confirming understanding of the specific machine we are resurrecting, and (b) validating programs written for it, and furnishing a realistic training aid for potential demonstrators.

*Contact Arthur Rowles at <rowles01@globalnet.co.uk>.*

### **Bombe Rebuild Project**

*John Harper*

With two notable exceptions there is nothing to report this time other than that in the near future we will have a major milestone to celebrate. There are 36 Letchworth Enigmas needed to complete the front of the machine: 118 commutators are required for this. All these commutators are now complete and will soon be built up into the last three Enigmas. Currently 33 are physically fitted and 29 are cabled up and fully tested. When the last three Enigmas are complete we estimate about another three days work will see the front of the machine complete. We are therefore on track to have the external aspects of the machine complete in time for the beginning of June events. The external covers are unlikely to be fitted at this time but perhaps that is just as well because this would make the interesting parts of the machine less visible.

Other items as previously reported have all made good progress and will be reported upon in more detail when we have more finished items to display.

The two major steps forward are in the area of Selenium Rectifiers and Special Resistors.

Acquiring Selenium Rectifiers to our original specifications has been a problem. For some years we put out calls for help but all that turned up was the odd single item, often not quite the correct type. We then decided to see if we could get them made. After a great deal of searching we were lucky and located a firm in the States which still made rectifiers. They came up with an item which is very close both physically and electrically to the originals. Better still, they gave us a below cost price and delivered in a few short weeks. We have now fitted 104 of these into the machine: they await their wiring.

Another item we had a great deal of trouble identifying was a wirewound non-inductive type of resistor. Again 104 are required. The problem was that we had been looking in the wrong place. We assumed that they would come, like the other wirewound resistors, from punched card equipment. It was a chance remark that led to us into realising that they were in fact GPO resistors. We now have two thirds of the number required and are not expecting any problems locating the rest. Most are the incorrect value and will have to be rewound. This led to another problem. Insulated resistance wire is virtually unobtainable but with a great deal of searching we found a firm with enough stock available, again in the States. This wire is currently on its way across the Atlantic. The interesting thing about this is that the supplier has told us that there will be no more of this wire manufactured unless somebody orders a quantity about 100-fold greater than our needs!

As I mentioned earlier, my next report should announce that the front of the machine is complete. This also means that the 12 miles of wiring will all have been completed, fitted and tested. Other items as mentioned in previous reports are progressing well. This leaves as the major task the various relays and associated circuitry to make and fit. After that we can start commissioning.

Our Web site can be accessed at [www.jharper.demon.co.uk/bombe1.htm](http://www.jharper.demon.co.uk/bombe1.htm).

---

## **Inmos and the Transputer (part 2)**

*Iann Barron*

---

This is the story of the UK's attempt to become a force in the emerging microprocessor industry of the 1970s, told by the man who was the driving force behind it. In this second part, which starts at the point when Inmos had just secured its second slice of Government funding, Iann Barron charts the company's decline and fall.

We then had to build a semiconductor plant in the UK. The Thatcher Government said "We're hands off. We don't involve ourselves with industry at all. We certainly wouldn't want to involve ourselves in any conflicts about where you locate the plant, but it had better be in South Wales"! So that's where it went.

We got the facility up, we got the process in, and then we hit trouble. We couldn't make the process work. We could make integrated circuits but they were unreliable and when we did tests on them we found they suffered from metal fatigue, and a small proportion didn't function correctly.

We identified a particular piece of process equipment which we didn't think worked very well. However the American company insisted that the fault must lie with the management of the operation – no problem with their process – it was the people who were running it in the UK.

Then, three or four months later, at Christmas 1982, we had a severe disaster. A semiconductor plant depends upon using copious quantities of very pure water and has enormous de-ionising facilities for the water to get to absolute, or very near absolute purity. I was rung up on Christmas Day with news that the water was anything but pure and was getting worse.

Even though we were now almost in full production, we had to close the whole plant down. It was not at all clear what had gone wrong. Eventually, after about three days, we identified that the problem lay with the resins which were used for the de-ionisation process – basically chemicals in huge 'tin cans'. The tin cans were about three metres in diameter and eight metres high. Something had gone wrong with them.

What we had to do – and this was all across the holiday period – was to physically dig out all this resin and get another lot trucked in and loaded. Virtually everybody in the company, including me, had shovels, and was standing in these great big tin cans shovelling resin.

The result was that the Americans said, "These Brits don't know what they're doing. Get rid of them!". The semiconductor facility was taken away and put under the control of the Americans who were deemed to understand these things.

What had actually happened, as we found out three months later, was that on Christmas Eve the engineers at the local reservoir decided to celebrate. They were supposed to stay on site, so what they did was to dump 100 times the standard level of chlorine into the water supply, then go off and have a Christmas party. That chlorine totally ruined our semiconductor plant.

We had many further problems. One was with money, because the pound collapsed and so my original £50 million had halved in value by the time we got around to spending the second half of it. So we ran out of money. We then had another highly protracted and visible round of arguments with the Government before we managed to extract another £15 million which was necessary to compensate for the fact that the second £25 million had been delayed.

We had just one staunch ally in the Thatcher Government – Keith Joseph. He consistently supported Inmos, consistently argued our case, and went to the Cabinet three times to get this money. While he was doing this he was being slated in the UK press for his lack of support for Inmos. A sad thing.

We eventually got our money. We got our process up. We got our SRAM out and Inmos became very successful. We were making a lot of money, and we had very good prospects. But Mrs Thatcher decided she didn't want Inmos, so it was politically unacceptable to recognise that we had been successful.

We suffered a succession of press and parliamentary statements about how unsatisfactory Inmos was. We also had the requirement placed upon us as management that we try to buy the company. I can tell you that it is extremely difficult to find a buyer for a company when your major shareholder is the Government and is saying that you're no good.

Our American partners did actually find an American company, which offered to take Inmos off the Government's hands for a fee of £15 million. Incredibly, this was agreed. So I went to the then chairman of Inmos and said, "Look this is totally wrong. We have a company that is doing extremely well and we have a deal which is being stitched up between an American company and our American partners to take it off the UK Government. For this they want to be paid – and then they are proposing to make a lot of money out of the US company."

My chairman said "Yes, I understand all that but I can't do anything about it. If you want to do something I will treat this conversation as privileged". So I went to see Mrs Thatcher and I told her that if the deal went through, in six months there would be so much mud on the face of the Tory Government that it would do them no good. That stopped it!

Back at the ranch, we were still trying to develop the transputer. I had the problem that I was trying both to direct the technical development of the transputer and to cope with my American partners and an unsatisfactory political environment. It was not a very pleasant time in many ways.

Going back to the transputer itself, we had two technical gurus, David May and Robert Milne. They were just like chalk and cheese – completely different in their approaches.

Robert was substantially more clever than David, and David's quite clever enough. When Robert wrote his PhD thesis in Edinburgh his supervisor Robin Milner, who was not lacking academically, said he couldn't understand a word of, it but seeing that Robert had written it, it must be right.

Unfortunately Robert did have the problem that he could not communicate very well with mere mortals. I spent my life trying to understand what it was he was saying. So we had a high level conflict between two people who were both very clear on their views on parallelism, one of whom had a simple vision and one of whom had a far more complicated vision.

Eventually I had to decide between the two. I had no hesitation, because David's position was basically where I had been some time before, but it was very painful at the time. Robert contributed a great deal to the underlying ideas and architecture of the transputer and he deserves recognition for that. Another thing he contributed, ironically, was the name Occam for our programming language.

In terms of process technology we had Peter Cavill from Marconi and Jonathan Edwards from Plessey. The depth of knowledge in the UK on process technology fundamentals was far better than in the United States, and most of the process ideas which were eventually developed came out of the UK company's microprocessor activities.

The CAD system and workstations were developed by Colin Whitby-Strevens and the actual detailed implementation and layout of the transputer was done by a team under Miles Chesney.

We made an experimental microprocessor to prove the technique, to prove the design system, and to prove the process. We had to test a lot of the novel logical techniques that were being used inside. We made a cut-down version of the final transputer, without the on-chip memory and many of the more complicated instructions that relate to processes. We built a second chip to go with it which enabled us to take the outputs and display them; this drove a CRT display.

It all worked! Just like that! Bingo! I think we had one minor revision, but essentially we had produced a complex microprocessor which worked first time.

Essentially a transputer allows a set of programs or processes which are independent of one another to run concurrently, so that the processor is shared between a set of processes. The processors communicate with one another by a simple channel which passes messages and provides synchronisation.

We built the instruction set hardware so that what the programmer saw was effectively a set of parallel processes. The microprocessor implemented them and scheduled them. It did the scheduling in an optimally efficient way in terms of the number of times one swaps from one process to another. Again it all worked!

The trick was that we made sure that the communication between the processes that were running on a transputer was identical from a programmer point of view, whether it was done within one transputer or between multiple transputers. This meant it was invisible to the programmer whether he was using one transputer or several. The program was the same; it was just compiled in a different way depending how many transputers were being used.

Of course there was a performance overhead, and we put a lot of effort into ensuring that the overhead for going from one process to another was as small as possible. It was equivalent to something like three basic instructions, so we could have lightweight processes running very short programs with a negligible overhead. Some transputer applications were written with processes as small as 10 instructions.

The other thing we did was to create the Occam programming language, which directly mirrored the structure within the transputer. Occam was not a one-for-one assembly language; it was a high-level language built with predictability in terms of the length of time taken to execute any bit of compiled program.

All that was great. Everything worked just as the theory said it should. There was just one thing that went wrong – and that was the compilers.

It was quite clear that we were moving to a situation where we were going to have a programming language for our own internal purposes that effectively replaced an assembly language, but we expected that customers would want to program transputers in high-level languages.

That was a fundamental strategic goal from the beginning. It was obvious that people could not go on programming microprocessors in assembly language. There had to be a change.

We thought that if we had a chip which was designed properly to support high level programming languages, we would then be in a very strong position when that change occurred, compared to our competitor companies who were locked into architectures designed for hand-coding, because for them it would be extremely difficult to exploit compilers. We were right about that.

Unfortunately we had a problem – money! We had no money to develop the compilers we needed. Our funds had all been used up in the US. We tried to get away with using a cheap untried UK solution in which we were going to get three compilers generated from one basic piece of code – and it didn't work!

So by the time the transputer came to market everything was there, everything was right, except that we didn't have the compilers necessary to exploit the capability we had created. We never did get those compilers.

Reverting to the original business plan, the proposition that went to the NEB was to make an investment in a semiconductor company, to get it up and running, and establish a viable semiconductor business in the UK. Then the intention was to sell it off, doing away with the need for the NEB by refinancing Inmos in the UK as a UK company.

The plan said that this would happen in 1984. I had forecast that in that year Inmos would have a turnover of £150 million – which it did – that it would have a profitability of £12 million – I think we actually made £15 million – that the NEB would exit during the year, and that the company would be valued at £125 million. That is exactly what happened.

I bet nobody's ever produced a business plan and got it as right as that. I have to admit that if you look at the forecast for any one of the preceding years, they were totally and utterly wrong. But 1984 I got exactly right!

The Government was desperate for us to be privatised in some way, but they had no takers. I devised something I called the white pawn strategy. You've heard of white knights – white pawns are the poor man's version. There was no company in the UK that looked like a white knight, but I did think there were perhaps half a dozen companies that we might persuade to invest in Inmos to provide an underlying resource to support their businesses.

One of these was Thorn EMI. Come the middle of 1984, the then Chairman of Thorn EMI, Peter Laister, started to get ambitious and wanted to acquire something significant. He bid for British Aerospace, failed, and attracted a lot of criticism. One bright young man in his private office said to him "Why don't you buy Inmos instead?", and so he did, in less than a week.

He paid £125 million for Inmos. It was quite sudden, totally unexpected and, from my point of view, very surprising. I was called up one day and told "There's this company that wants to buy you". I went to various meetings, the last of them at the Department of Industry offices in Victoria Street. It started about midday and went on until midnight at which point the deal was done.

I was excluded from most of the meeting. I had no direct involvement, and I was told, "Just go down the corridor and find somewhere to sit". So I did. I sat down in an office and after about 10 minutes I started to look round and realised it was rather a grand office. I looked a bit more and realised there were papers on the desk. I was actually sitting at the Minister's desk. That's security for you!

Eventually I was called back into the meeting and was told the deal had been done. Peter Laister produced a transputer chip from his pocket, waved it at the assembled audience, which consisted of one person from the NEB, one from the Government and 50 lawyers and said, "Look, I've paid £125 million for this".

There was one flaw. The business plan said that in 1984 not only would the company be sold, but that it would be refinanced, because we knew at that stage we would want money put in to expand our business further. But there was no refinancing, and no money went into Inmos. The £125 million went straight to the Treasury.

So from that moment on we had no cash to do anything. We couldn't advance the technology in any way. We couldn't develop our memory products. It was not really Thorn's fault. They shouldn't have bought Inmos – it was a really stupid thing to do, and Peter Laister was ousted from that company within three months. Thorn did try their best to look after us. It was just that they didn't have the kind of cash that was necessary to fund and run a large-scale semiconductor company.

We had been doing very well. Our static RAM products, in spite of the technical problems, had knocked Intel out of the market. We had something like a 60% share of the world's market for them and our only competitors were Japanese companies. We had dynamic RAMs of high performance, and we had transputers just about coming to market. So we were in a very good technical position. We just didn't have the money to exploit it.

There was one other unfortunate occurrence at this stage. Almost immediately after selling the company, it started to become apparent that some of our customers were having failures in their static RAM memory chips. This problem escalated, and after about six months it became clear that we had a serious reliability problem.

The problem was due to a faulty piece of equipment. Inmos was not the only company to suffer at that time and two or three semiconductor companies actually went under as a result. Which piece of equipment was it? Yes, you've guessed it, it was the problem we had identified in the UK factory when we started to manufacture the static RAM in 1982. At that time, not only had the US taken over control and management of the UK factory, but they had also transferred production of the static RAM back to the US – it was a highly profitable product, and the company making the static RAM

automatically had the better balance sheet. Heads rolled, and much of the US management was dismissed.

So we had to cut back the company and lost half our transputer design team. We tried to retrench in an appropriate management fashion and to create something valuable out of it, so we actually hived off some of our best people into a new company which then went into business making parallel computers.

The story got worse. Eventually I had the ironic situation where I was in the US running the US company and I had to sell it off. It was actually sold to Seymour Cray, whom I had spent quite a long time trying to convince of the value of the transputer.

It wasn't all doom and gloom because we had, as I have said, good technology. We had developed fundamental patents for static RAMs, for dynamic RAMs and for microprocessors, to the extent that every static RAM or dynamic RAM made involved a significant patent from Inmos. IBM was so concerned about our patent position in microprocessors they gave us rights to all their patents in return for rights in our patents.

Eventually dear old Thorn made something like £150 million to £200 million out of the Inmos patents, which is hardly ever mentioned. They got their money back all right.

There were two other money spinners. Our dynamic RAM technology was very good, and we were able to licence it to a Japanese company and set up an automated manufacturing facility in Japan. The other product was the second test chip we had made for the transputer. We persuaded IBM to use this as their next generation graphic chip, and it became embedded in the PC as the SVGA standard. Some of the oddities in the present Windows operating system stem directly from the way this chip was designed to work with the transputer.

But from the point of view of the transputer we did not have the resources to do what was necessary. We did do one clever thing which was to add a floating-point unit to the existing microprocessor. We exploited the improved process technology to be able to put on board a floating-point unit and again we used formal techniques. We were able to prove what we were doing and we were able to design a floating-point unit better than anything else on the market by a very long chalk. That gave us another couple of years lease of life.

The end came in 1989 when Inmos was sold to SGS Thompson. SGS Thompson was essentially a semiconductor company set up by the French and Italian Governments to do what Inmos had planned to do in the UK.

The Managing Director, Pasquale Pistorio, was very enthusiastic about the transputer. He thought it was an extremely good product and wanted to revive it. He was prepared to put a substantial amount of money in this. But there was a fundamental problem. No work had been done on the transputer in terms of implementation and design since about 1985, and this was 1989.

There were two routes that we could go down. One was to take the existing design and just re-implement it again in the latest process technology, taking advantage of various tricks and getting more memory on the chip and things like that. The most favourable view was that this could be done by 1991; my personal belief was that it would take at least a year or two years longer and would not be competitive.

The alternative was to start from scratch and design a new transputer, keeping the concept and making sure it was software-compatible and connectable and things like that, but changing the instruction set and being quite radical.

The quite radical thing was to put four transputers onto a chip. You could say that we were doing much the same thing as the very long instruction word microprocessors. We were going to produce a system with a full 64-bit internal bus with 32-bit addressing, four address units, four arithmetic units, and two 64-bit floating-point units. We were going to put a 50 Mbps link between the transputers and everything was absolutely divine. The trouble is that this would take a long time to do, and it wouldn't have been available until 1992 or 1993.

So there was no real way for the company to survive. A decision was made to go the route of re-implementing the existing design in the latest technology, and it was a failure as expected. That was the T9000 – it never really saw the light of day. That was the last of the transputer and very recently ST Microelectronics, as SGS Thompson has become, announced that there will be no more transputers, which is sad.

If we look back now, it is clear that Inmos was a glorious failure. It had enormous technical successes. The static RAM technology design is the way that the world now designs static RAM. The same thing is true for dynamic RAM technology. We had the transputer which was a unique design, and we

had, as I have said, this very strong patent base. Our problem was the corporate failures.

The underlying issue was the control of the company. There were conflicts between the US and UK, which led to extreme mistrust. We had an antagonistic shareholder. Originally there were difficulties within the Labour Government, because there was one minister who objected – Tony Benn. Later our antagonistic shareholder was Margaret Thatcher.

I had a number of discussions with her aides about the problem. Mrs Thatcher continually said her aim was to support sunrise industries, and that she wanted to compete with the Japanese, and Inmos, of course, was doing all of those things. I was told that she had a very simplistic view towards life and liked to get things down to one basic idea. Her one basic idea about Inmos could not be changed, and that was that Inmos had been started by Tony Benn. Madam didn't like him very much.

At a management level we had the quality failures. There was also the underlying issue, which hits UK companies time and time again, that you can get the money to start, but when you get market success you really need more and you cannot raise it. What went wrong with Inmos more than anything else was the failure to get that second round of finance.

I think you will find that microprocessors in the future look more and more like transputers. You won't see it from the outside but it will be embedded in the way that they actually do their operations.

For all the problems, though, it was great fun, and I'm glad I did it.

*Editor's note; this is an edited version of part of the talk given by the author to the Society at the Science Museum on 19 November 1998. The first part, describing how Inmos came into being and the thinking behind the Transputer, appeared in issue 32.*

*Iann Barron is at <iann@mynchen.demon.co.uk>.*

### **Editorial contact details**

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

---

## **A Management Perspective on Orion**

*Peter Hall*

---

One of the main protagonists in the Orion story, Peter Hall, was unable to attend the Society's seminar on Orion 2 last October due to illness. He has since watched the video tape of the event, which has prompted the following comments.

I make no apology for my decisions which led to the real problems with Orion 1, but I think it might be interesting to put on record how they came about and make some comments about the consequences.

I took over the Computer Department in 1958, having been given the job because the then Manager Brian Pollard had left and a few years previously an old wartime friend and colleague Bill Elliott who had been running the southern laboratories – Lily Hill – had also resigned. Bill had taken a few key people with him to IBM!

I can only assume that I was given the job because I had rescued another ailing department – it certainly was not because I knew anything about computers – I knew nothing. Perhaps I should add here that on taking over the Department I got no briefing from any of the senior Ferranti management – except "Put it right"! I had no briefing from Brian Pollard: he had gone before I arrived.

Ted Brauholtz's reports of meetings with Sebastian Ferranti and with Brian Pollard concerned with a justifiable warning about "neurons" were news to me. I heard of them first when I spoke to Ted a few weeks before the seminar. If only he had been able to get hold of me in 1958 the story might have been different.

The southern team at Lily Hill had been responsible for the development of Perseus. Delivery was late and the machine for the Old Mutual in Cape Town was very late indeed. In July 1959 I had to go to South Africa with Peter Hunt to placate a very irate customer who had written a vitriolic letter to Vincent Ferranti. This was not an easy trip, and while I judged that the Lily Hill lot was a competent gang, from my standpoint they had certainly made my first year tiresome! Perhaps the previous management had given them a job too big for the resources made available.

The team at Lily Hill had, it seemed, been left after Bill Elliott had gone without a clear management structure. I soon put Peter Dorey in charge. He later became one of Ferranti's senior and most successful managers and a director of the company.

Peter Dorey expressed no desire to take on another major project, and I gave no thought to getting them in on the Orion act. With hindsight this was clearly a mistake.

The West Gorton situation was different. Mercury and Pegasus were in production, seemingly reasonably on time and not causing any serious headaches. Perhaps here I should correct a statement made at the seminar regarding the University machines.

The university made what used to be called "bread boards"; while West Gorton was responsible for turning the circuits and the logical design into an engineered product, and for developing enhancements. Sales were going OK and the only problem seemed to be that we were making too much profit on Pegasus and we would have to give some money back to NRDC!

But there was a flaming row going on about what computer we should develop next. So it wasn't just the war between north and south I had to cope with but between two factions in West Gorton – the university protagonists and the "in house" people, who rarely seemed to talk to each other.

Fortunately I knew the university people well. Freddy Williams and Tom Kilburn were at the same wartime Radar Research Establishment as me (TRE), and I knew of their work. They were an exceptional team at TRE and I knew that their work at the university would be first class. Tom's proposal for a machine sounded exciting and would mean that we would establish a world lead in high speed computing. Could we do it? I judged that together we could.

On the other hand was the internal faction led by a man with a proven reputation (in imaginative thinking and in circuit design from Cambridge) – Gordon Scarrott. Gordon later proved his worth leading the teams on CAFS and DAP. Their proposal was for a more modest machine geared to the commercial market and to the cheaper end of the scientific market.

The circuits Gordon proposed, and had christened "neurons", had been tested in an experimental machine they called Newt – they seemed to work well. Later Newt was re-engineered and was sold as Sirius. Sirius was a success – again proving the circuits. The circuits were attractive because they were extremely powerful logically, and also from a production standpoint because many fewer types would be required than if we had used competing circuits. How wrong we were!

I made up my mind that we would go for both machines. This put far too much strain on all our resources – particularly when both programmes ran late; cash got more than a bit short! The neurons became our biggest problem.

It was not just that current-driven logic was difficult to commission; the circuit tolerances were just not good enough when used in such large numbers, and on such a physically large machine. We could have used "griblons" (as in Orion 2) from the word go! But nobody suggested we should.

I should have had a closer look at the neurons. Remember, though, that such was the infighting that nobody could be trusted to give an unbiased view. I wish Ted had got through to me – and I wish I had had "more hindsight sooner"!

I would like to make a comment regarding the relationship between Orion 1 and Orion 2. I felt (but I could not hear half of the soundtrack so I may be mistaken) that more acknowledgement should have been given in the seminar to the fact that the Orion 2 project started from a computer design that was known to work – Orion 1. The change in circuits meant a significant logical redesign, but essentially the job was only to re-engineer Orion with "Griblons" instead of "Neurons" – nothing else. With an architecture that was known to work and proven software – OMP in particular – it seemed to me a straightforward job.

In these circumstances it was possible to agree a very tough development programme. Indeed if it was not done quickly it would not meet our urgent need as an Orion 1 replacement. Although the Prudential requested some additional facilities the team did a fantastic job in meeting the delivery without incurring penalties. It is interesting to note that Orion 1 took roughly five years to develop compared to about three years for Orion 2.

A small point, but perhaps important, is that Nebula was not an Orion 2 project – we did it for Orion 1. It should not just be bracketed with Orion 2. We did hope of course that Nebula would be carried forward into later machines. I hope the contribution made to the Nebula programme by Peter Hunt was properly recognised. I believe Peter was a very outstanding software project manager.

There was a suggestion at the seminar that Ferranti should have done more market research before embarking on development projects. This is a difficult one. Rely on market research and you develop a Perseus, sell only two, and lose a fortune! If you just get on with the job, and stop thinking (as Freddie Williams told Tom Watson of IBM ) then you make an Atlas, and, you may be surprised to hear, make a little money! Seriously, there must be both "technology push" and an understanding of the market – "market pull". (But don't pretend customers know what they want).

Finally I should like to echo what Ted said in relation to our memories. We are looking back over 40 years and memories are fallible. Some things can be checked from various archives of course, but most of my comments are from my memory of those events long ago! I know I get things wrong (I have proved it to myself in relation to some wartime work).

*Contact Peter Hall at <[peterdhall@hotmail.com](mailto:peterdhall@hotmail.com)>.*

### **North West Group contact details**

*Chairman* Tom Hinchliffe: Tel: 01663 765040.

Email: [tom.h@dial.pipex.com](mailto:tom.h@dial.pipex.com)

*Secretary* William Gunn Tel: 01663 764997.

Email: [william.gunn@ntlworld.com](mailto:william.gunn@ntlworld.com)

*Science & Industry Museum representative* Jenny Wetton, Museum of Science & Industry, Liverpool Road, Castlefield, Manchester M3 4JP.

Tel: 0161 832 2244. Email: [j.wetton@msim.org.uk](mailto:j.wetton@msim.org.uk)

---

## **Before Moore's Law**

*Brian Wichmann*

---

Gordon Moore made some statements about the development of integrated circuits in 1965. However, it was not until 1975 that the phrase 'Moore's Law' was used. Even then, it only caught the public imagination in the 1980s. If you make a Web search for Moore's Law and see a paper by Ilkka Tuomi, then the history is revealed, but here is a short summary.

Moore's first statement concerned the number of components on a semiconductor that provided the minimum cost. At any given time, increasing the number of components beyond this threshold reduced the yield, making it less economic. The number defining this threshold seemed to be doubling every year.

This prediction was reasonably sound from 1959 to 1975. But the real driver to increase the number of components was the ability to provide functionality to a large market, via the general-purpose microprocessor.

In 1979, Moore's key message concerned the number of components on the largest chips, which gave a three year doubling rate over the period 1959-1975.

There are some fundamental problems with an analysis in terms of components per chip. As the technology improves, so does the packing density. Also, some chips, such as memory chips, are much simpler than others, such as a microprocessor.

The advent of memory and cpu chips changed the market completely. Intel's own figures for microprocessors over the period 1971-2001 show initially a doubling of the transistor count over 22 months, then over 33 months, and then in the last eight years over 54 months. The rate is clearly slowing down.

This analysis does not address the real point, which is the performance supplied to the user. Unfortunately, the industry tends to advertise performance in terms of the clock rate, while the user really wants to know how long a specific activity will take. For scientific computing a compromise would be the number of floating point operations that can be executed per second. On this basis, the Intel figures for the period 1971-1995 give a doubling rate of about three years.

## ALGOL performance summary

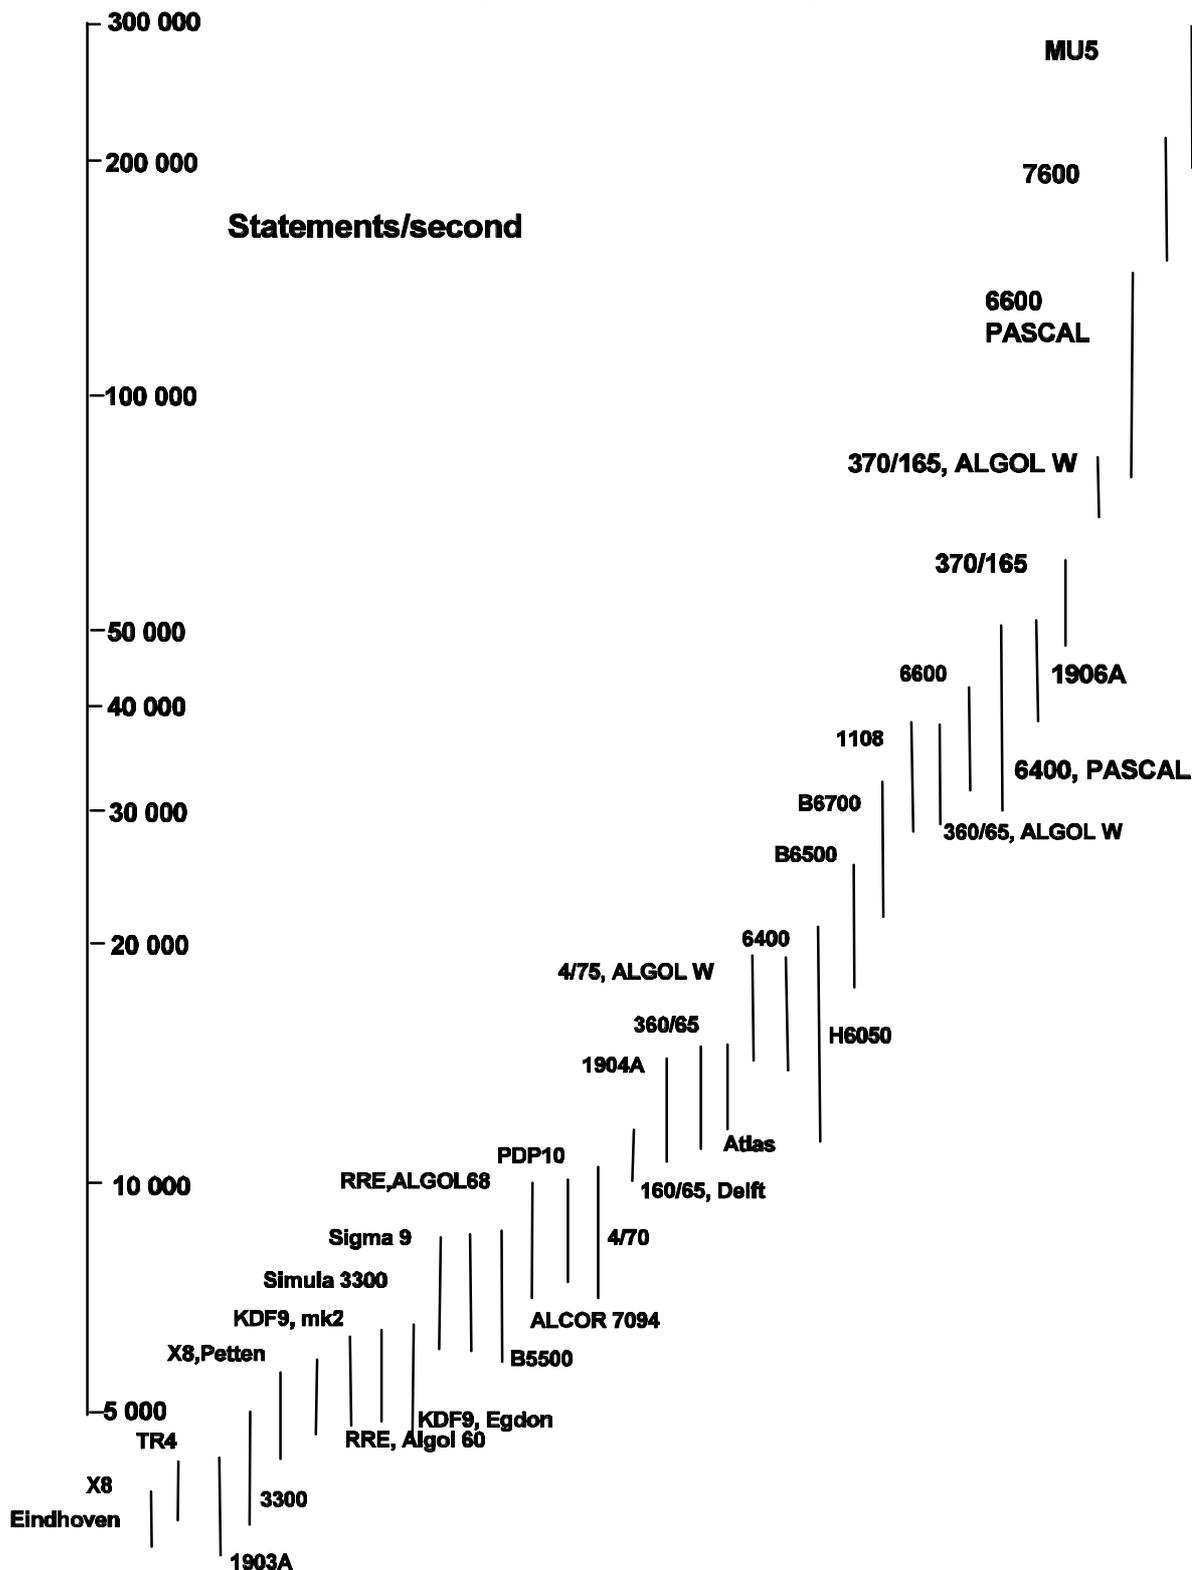


Figure 1: Statement time analysis for 34 machines

Measuring actual performance is difficult, since the relative speed of two machines depends critically on the task undertaken. When Intel published technical details of the 386 processor, it used the Whetstone benchmark which I wrote with Harold Curnow. However, Intel's coding (in C) was not a correct transcription of the original program. It is very hard to avoid such problems.

My interest here is in the speed of computers during the period 1960-1977. From 1965 to 1977, I spent some time at the National Physical Laboratory (NPL) studying the speed of computers and also the effectiveness of Algol compilers. Over this period, it was possible to move programs in Fortran or Algol 60 to a variety of machines and measure their relative speed. For Algol 60, it was possible to measure the times for individual statements, partly due to the absence of any effective optimization!

### **Algol performance**

My first study collected the execution times of some basic statements in Algol 60. These were all given in microseconds – hardly an appropriate unit today. An analysis was then undertaken on the basis of a simple model that the time for any statement should be approximately the product of a factor depending on the complexity of the statement and another factor inversely proportional to the speed of the machine. The analysis was then presented by taking Atlas as the machine of unit performance.

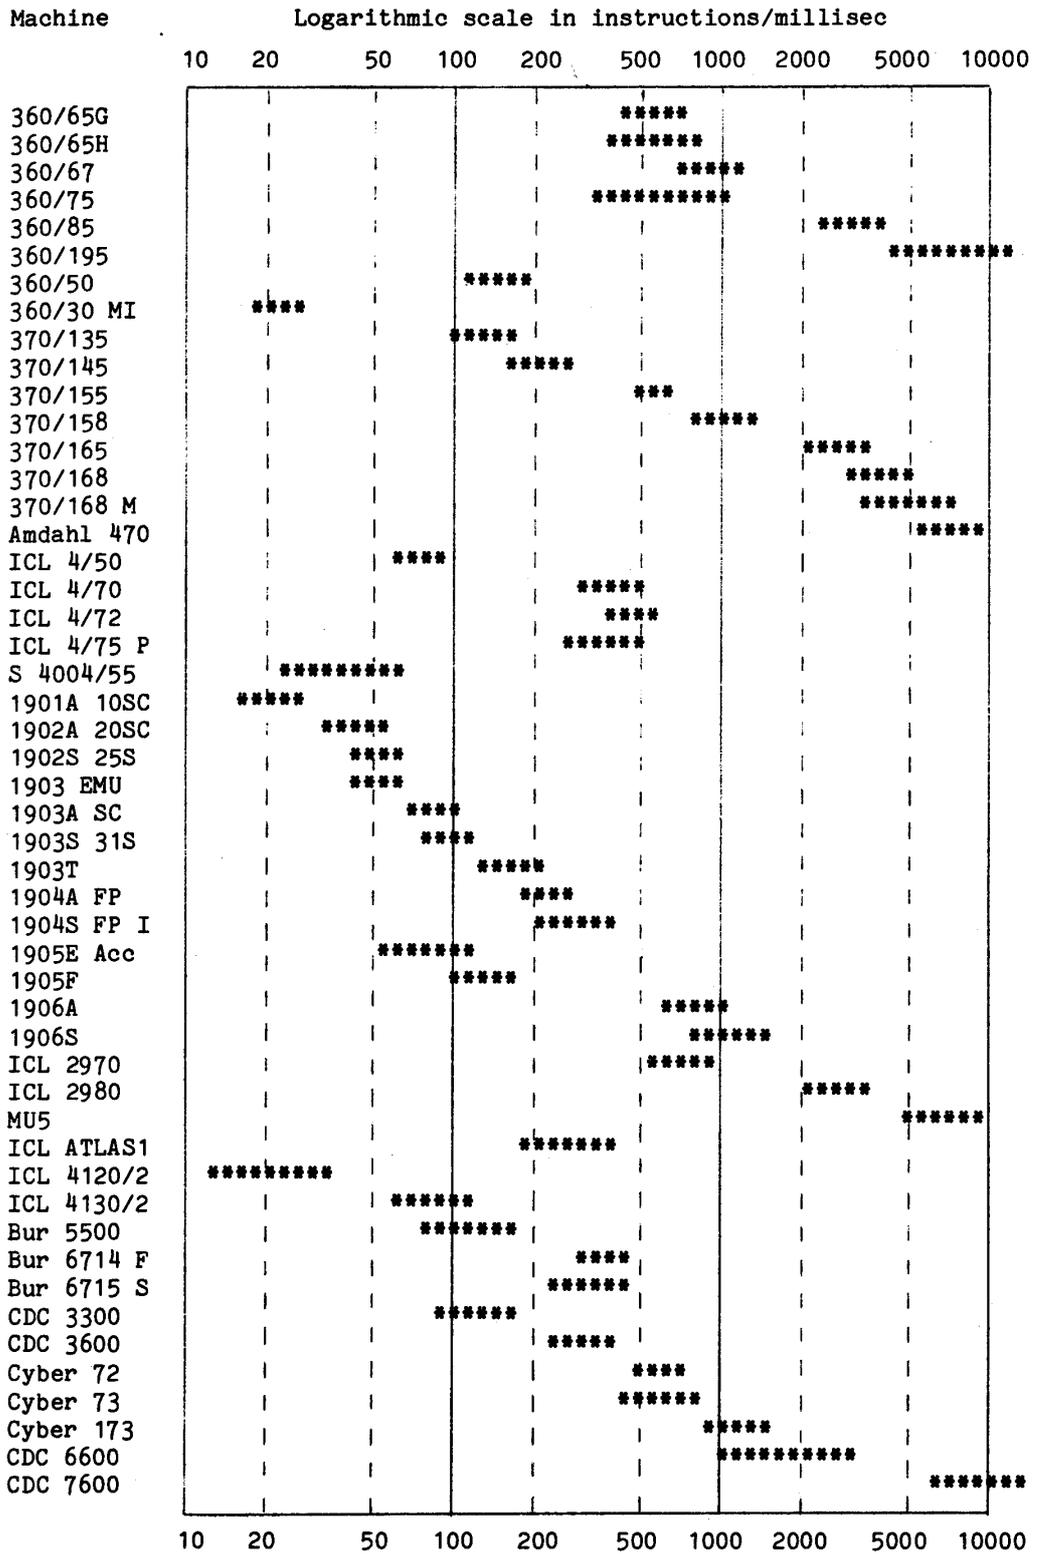
The use of Atlas in this regard was even enshrined in a Ministerial statement which required Government to go single tender to ICL for machines more powerful than Atlas.

The analysis of these statement times indicated which statements were implemented well and which were poor. Another study indicated the main reasons for the differences, at least for five compilers (again, including Atlas). The last report collecting these statement times included a graph of the performance of 34 machines. This is shown in Figure 1.

### **Machine performance**

These performance analyses failed to answer some basic questions, such as what *should* be the performance of a particular computer? It was clear that some Algol 60 compilers did show some machines in a good light and that other languages, particularly Fortran, would provide very different results.

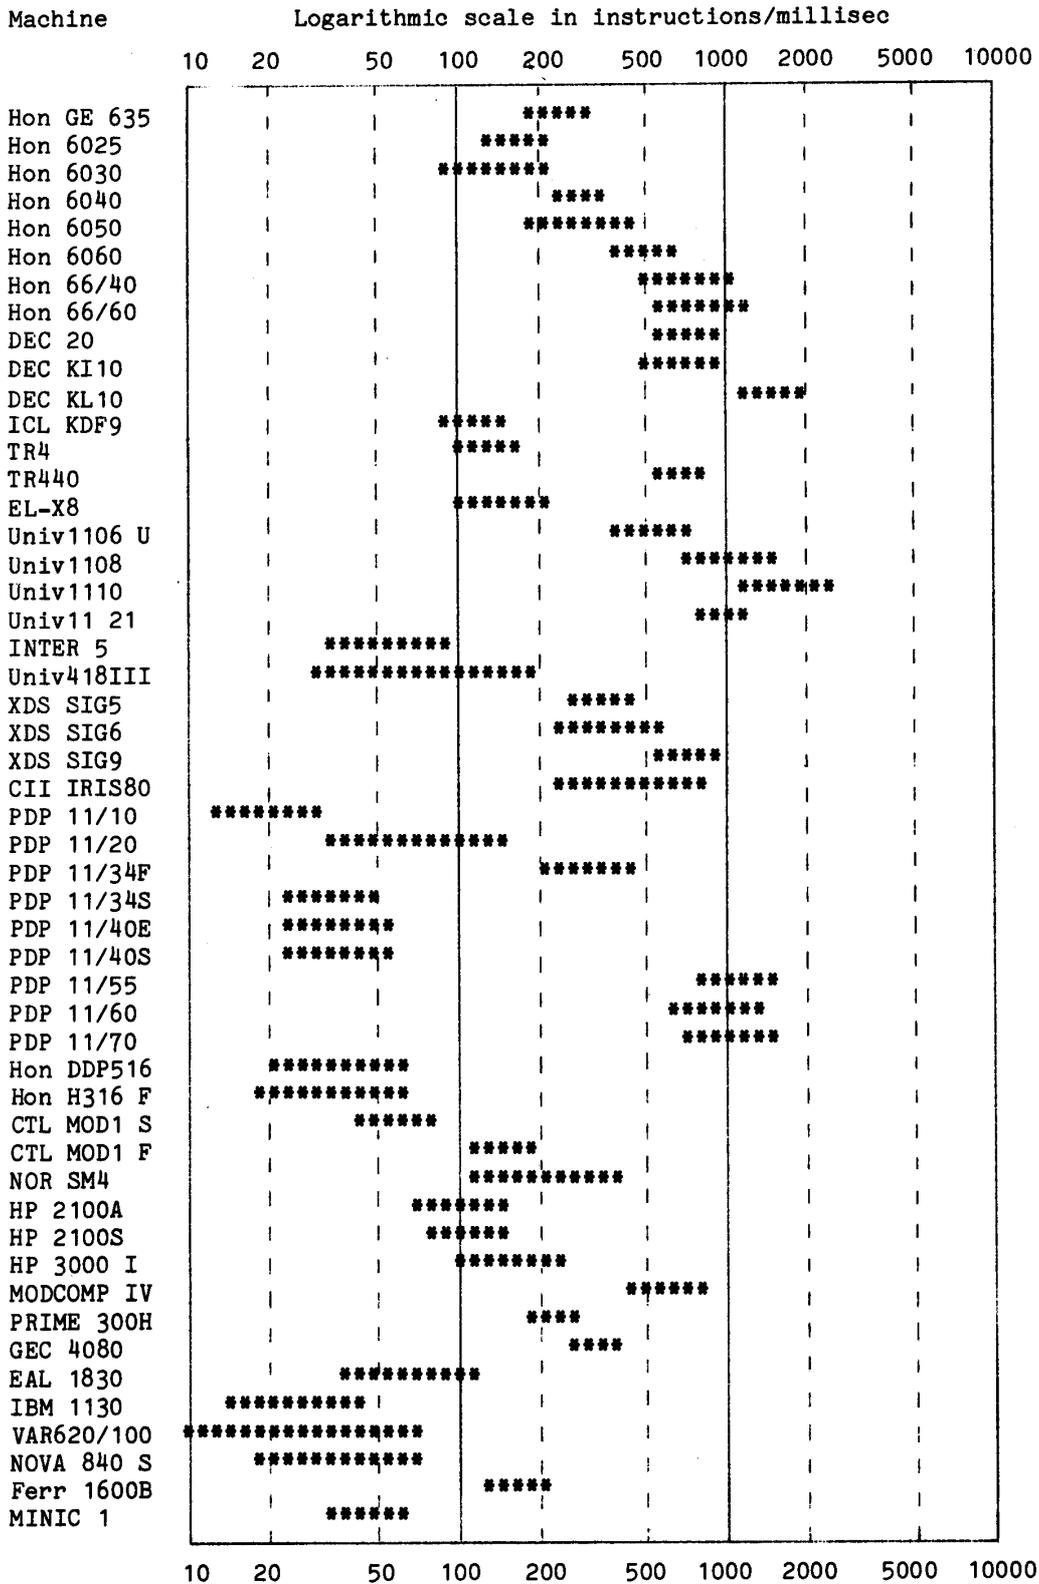
Performance Summary of all the Computers



CROWN COPYRIGHT

Figure 2: Summary, Part 1

Performance Summary of all the Computers



CROWN COPYRIGHT

Figure 3: Summary, Part 2

A small database was constructed and a substantial amount of information was provided by the Central Computer Agency (CCA) in addition to that from NPL. The CCA computed the Gibson Mix and ADP Mix for machines procured by Government. Also, as part of the acceptance test, they ran some benchmark programs and timed them. A research report documenting the results about this created a lot of interest, so it was decided to produce a 'commercial' report using the same analysis technique and incorporating further information from CCA.

Clifford Nott spent much time and effort producing a *Guide to the Processing Speed of Computers* which would do credit to NPL and justify its market price (£50). When 200 copies in a flashy red ring file were ready, I produced a press release to mark the occasion. This was duly presented to the then Director of NPL, Dr Paul Dean. Unfortunately, he decided the material was too 'commercial' being perhaps like a *Which?* report. I suspect I have the sole remaining copy of the (non-)publication. The crucial figures giving the performance of the machines appear in Figures 2 and 3.

### **Postscript**

The data on performance which relates to Moore's Law is really only applicable when microprocessors became the main computing engines for large parts of the industry. In contrast, the data that I collected in 1977 related to large mainframe computers. However, the two are clearly related and indeed, the benchmark programs can easily be run on modern microprocessor-based systems to give a direct comparison.

The long-term success of computing is amazing. When I retired in 1999, I calculated that my home computer (an iMac) was 13,000,000 times more cost-effective than the KDF9 which I used when I started work in 1964. This represents a doubling of cost-performance roughly every 18 months.

Unfortunately, my data lacks any indication of when the computers first came into operation. If this could be added, then an analysis of the increase in computing power over the period 1966-1977 should be possible. I hope to work with Professor Bob Hopgood to add this information so that the extrapolation of Moore's Law can be determined over the period 1960-1977. Hence I would be delighted to know the relevant dates for any of the machines listed.

*Editor's note: Contact the author at <Brian.Wichmann@bcs.org.uk>.*

## **Obituary: John Gosden**

---

*David Caminer*

---

John Gosden, software pioneer, died in New York on 18 December 2003 at the age of 73.

Gosden joined J Lyons as a trainee programmer in 1953. At that time its small Leo team was engaged on the final trials of the payroll programme for the Cadby Hall bakeries. The Leo system of the time was equipped with only the most rudimentary systems software: the programmer had to do it all. By the time that Gosden left six years later he had sketched a full body of systems software for the latest Leo system, based on his rich experience in implementing programs of all kinds and on his study of the most advanced work elsewhere.

A key part of Gosden's work was in working alongside John Pinkerton, Leo's Chief Engineer, to ensure that Leo III incorporated all the facilities required to meet business application needs. This system was able to run several programs at the same time and was microprogrammed to enable users to devise new macroinstructions to implement frequently used sequences. Gosden also devised the high level language Cleo for the system.

On joining the Auerbach Corporation in the United States in 1960, Gosden raised the Standard EDP Reports to new levels and then managed the construction of a multi-computer operating system for the US Navy. Later, at the Mitre Corporation, he was responsible for planning the database requirements for the Joint Chiefs of Staff, and then at the Equitable Life Assurance Society of New York he was VP responsible for major projects and policy. Among other activities over the years he chaired a committee advising on White House support systems. Well into his retirement he advised the Museum of Modern Art and the Cornell Medical Centre on their database systems.

Notwithstanding his crowded US career, John Gosden remained very close to his Leo roots. He was Associate Editor of *LEO, the Incredible Story of the World's First Business Computer* in 1987 and contributed a chapter entitled "Toward Systems Software". In 2001, although already in failing health he crossed the Atlantic to contribute to the fiftieth anniversary celebrations of the first business computer application in London's Guildhall.

---

## Letters to the Editor

---

Dear Editor,

I was interested by Iann Barron's article on Inmos and the Transputer in Issue Number 32 of *Resurrection*.

The DAP (Distributed Array Processor) preceded Inmos. I remember Iann visiting the ICL research labs in Stevenage to look at the Pilot DAP (which had 1024 1-bit Processing Elements) as part of a government inspection delegation, I think early in 1976. I remember Iann holding one of our boards with 16 PEs on it, and looking thoughtfully at it for some time. We had been advised not to engage in argument with the delegation, so there was little discussion. I would be interested if Iann remembers the visit, and if it had any influence on his thinking about the Transputer.

After considerable rivalry and many generations of both the SIMD DAP and the MIMD Transputer, they both more or less died. However, there is what might be described as a joint successor technology alive today. The DAP lineage ended with CPP (Cambridge Parallel Computing), which folded in March last year. The previous year I left CPP to join a startup, WorldScape Inc (<[www.wscapeinc.com](http://www.wscapeinc.com)>), that is involved in SIMD technology in association with ClearSpeed (<[www.clearspeed.com](http://www.clearspeed.com)>), a company (formerly known as PixelFusion) in Bristol that is an indirect descendant of Inmos. There is currently a 64-PE chip available which includes IEEE floating point units in every PE. There are traces of Transputer heritage in the architecture, but it is essentially SIMD technology from some of the Inmos MIMD veterans. A new beginning.

Returning to the early years. I first met Iann at a computer workshop in Grenoble in June 1973, where I presented the first paper on ideas that became the DAP, but I don't expect this caught Iann's attention. The biggest players in the early DAP scene were David Hunt and Peter Flanders in my Stevenage team, and Dennis Parkinson in ICL technical marketing. My manager at the start was John Iliffe.

One day I may write more on DAP history.

Stewart Reddaway

By email from <[Stewart.Reddaway@wscapeinc.com](mailto:Stewart.Reddaway@wscapeinc.com)>

17 February 2004

---

## Forthcoming Events

---

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

**13 May 2004** AGM, followed by London meeting: “The History of Health Computing in the UK”. Speaker Professor Bernard Richards.

**1 June 2004** Colossus Anniversary meeting, Science Museum.

Details are subject to change. Members wishing to attend any meeting are advised to check in the Diary section of the BCS 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 at 1730; tea is served from 1700.

Queries about London meetings should be addressed to Hamish Carmichael, and about Manchester meetings to William Gunn on 01663 764997 or at <william.gunn@ntlworld.com>.

---

## CCS Web Site Information

---

The Society has its own World Wide Web (WWW) site: it is located at <[www.bcs.org.uk/sg/ccs/](http://www.bcs.org.uk/sg/ccs/)>. This is in addition to the FTP site at <[ftp.cs.man.ac.uk/pub/CCS-Archive](ftp://ftp.cs.man.ac.uk/pub/CCS-Archive)> (please note that this latter URL is case-sensitive). Readers wishing to access past issues of *Resurrection* should use this FTP site. The Web site includes information about the SSEM project as well as selected papers from *Resurrection*. Readers can download files, including simulators for historic machines.

---

## 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* **Hamish Carmichael FBCS**, 63 Collingwood Avenue, Tolworth, Surbiton, Surrey KT5 9PU. Tel: 020 8337 3176. Email: hamishc@globalnet.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* **Dr Xerxes Mazda**, Manager of Collections Access, Science Museum, Exhibition Road, London SW7 2DD. Tel: 020 7942 1100. Email: x.mazda@nmsi.ac.uk

*Museum of Science & Industry in Manchester representative* **Jenny Wetton**, Museum of Science & Industry, Liverpool Road, Castlefield, Manchester M3 4JP. Tel: 0161 832 2244. Email: j.wetton@msim.org.uk

*National Archives representative* **Jeffrey Darlington**, Digital Preservation Department, National Archives, Kew, Richmond, Surrey TW9 4DU. Tel: 020 8392 5268. Email: jeffrey.darlington@nationalarchives.gov.uk

*Computer Museum at Bletchley Park representative* **Michelle Moore**, The Computer Museum, Bletchley Park, Bletchley MK3 6EB. Tel: 07748 981391. Email: michellejqmoore@hotmail.com

*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**, 10 The Vineyard, Bouldnor, Yarmouth, Isle of Wight PO41 0XE Tel: 01983 761399. 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, DEC Working Party* **Kevin Murrell**, 25 Comet Close, Ash Vale, Aldershot, Hants GU12 5SG. Tel: 01252 683503. Email: kevin@ps8.co.uk

*Chairman, S100 bus Working Party* **Robin Shirley**, 41 Guildford Park Avenue, Guildford, Surrey GU2 5NL. Tel: 01483 565220. Email: r.shirley@surrey.ac.uk

*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, Mil-DAP Working Party* **Brian M Russell CEng MIEE**, 5 Briarmere Walk, Chadderton, Oldham OL9 6SH. Tel: 0161 652 6475. Email: bmrussell@iee.org

*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

---

## Committee of the Society (contd)

---

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

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

**Dr David Anderson**, The Woozle, 1 Oatlands Road, Boorley Green, Hants SO32 2DE. Tel: 0239284 6668. Email: cdpa@btinternet.com

**Peter Barnes FBCS**, 10 The Broadway, Gustard Wood, Herts AL4 8LN. Tel: 01438 832906. Email: barnes@peterbarnes.freeserve.co.uk

**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: mck@dcs.warwick.ac.uk

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

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

**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 Hamish Carmichael 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