

My involvement with Atlas: reminiscences of events, and how Atlas shaped my later career.

Firstly I should say that I feel a bit of a humbug being on the rostrum alongside such eminent contributors to the Atlas project, of which I was a very junior member. Because of this my contribution will be somewhat anecdotal, rather than deeply technical or historically significant. Secondly, I have to remind you that I will be talking about events that took place 50 years ago. So I should like to apologise in advance for any of the inevitable errors or omissions.

I joined the Atlas Supervisor team in the autumn of 1961, straight from Cambridge University, where I had just completed a PhD in maths and electronics. I had used EDSAC 2 extensively in my work, but as a user for a specific application, never from the point of view of the system as a whole. However in those days all programming was done in assembler, so you had to get very close to the structure of the machine, organise the layout of your data, your overlays etc. My work had also involved the use and maintenance of a special purpose digital computer, so I had some understanding of digital technology, including coping with unreliable systems and dealing with systems errors. I also had had experience of writing and running large programs, mainly for solving large sets of non-linear differential equations. These turned out to be very similar to the kind of programs many of our potential users wanted to run.

All of this would prove very useful and relevant to my work on Atlas and subsequent systems. I suppose that it was because of this background that I was taken on, since at that time that kind of experience was, I imagine, fairly rare. But I didn't realise at the time what an honour and a privilege it was to join this particular team. Not only were we working on what we hoped would be the world's fastest computer, which was exciting enough, but I was also joining a team of programmers of outstanding ability. In particular, it was a privilege to work under Dave Howarth who was a truly exceptional programmer. I didn't know at the time that his qualities had been recognised by the company when he was recruited to lead the Atlas Supervisor team. I was recently shown a report written some years later by Bernard Swann, who was Ferranti's Sales Manager in London at the time. In it he said:

“Dave Howarth was one of the most extraordinary workers in a field which has produced a number of men and women with remarkable mental and physical capacity... He could continue at a level of activity which would have exhausted most men.”

I say Amen to that.

We had to travelling up to Manchester to run programs on the machine, and we had to work at night, because the engineers were still working on the hardware during the day. We would run programs into the small hours, and then collapse back to our hotel rooms. Dave, however, would work right through the night, and then, on the train to London the next morning, would be working solidly on all the printouts, planning his next set of runs. And all he seemed to need to sustain him was a continuous stream of cigarettes!

It was always a matter of great regret to me that Dave's outstanding abilities were not more fully recognised later on in his time in ICT and ICL. I think this was because he was essentially a technical genius, and was not interested in playing political games or managing large departments.

To quote the Bernard Swann report again,

“One of the remarkable features of the Atlas project was the small number of staff compared with the large numbers employed by IBM on Stretch, and by the Bull Company on the Gamma 60 (which were comparable projects)... The former had some 300 graduates and the latter about 200 programmers. Atlas never had more than ten programmers on the Supervisor and about 15 working on compilers.”

I don't know what all the IBM and Bull staff did, and I don't know what Dave thought about having such a small team, but all I can say is that it seemed to more or less work, although there were delays towards the end of the project. But one of the advantages, to my mind, of having such a small team was that all of us had a good understanding of the overall picture, and worked together very closely, mostly very informally. I don't see how you could have achieved that with a much larger team.

I can remember most of the members of the supervisor team. Apart from Dave Howarth, there was Peter Jones, Mike Wyld, Mike Baylis, Jim Doughty, Peter Warne, myself, and possibly one or two others.

I don't remember having any formal induction into the team or the Supervisor, although no doubt I was given lots of documents to read. I was more or less told what my first task was, and was expected to get on with it, consulting others as necessary. The first task I was given

was to design and develop some test programmes to check the validity of the central routines of the Supervisor. This of course required me to have a very detailed understanding of the workings of these routines and how they interacted with the rest of the system.

The time I was beginning to work on this happened to coincide with an upsurge in marketing activity in relation to Atlas, and the sales team were often calling on Dave to go and give talks to potential customers about the Supervisor and the design of the computer more generally. I got the impression that Dave was very annoyed about this, partly because it took him away from his programming, and tended to delay the whole project because of the key role he had in every aspect. But it was also, I suspect, because Dave didn't actually like very much talking to outsiders about the project. Despite his brilliance, he was a fairly shy, unassuming person, who just wanted to get on with the job. And so I guess he looked around for someone to whom he could delegate this activity, and he picked on me. Fortunately, by then I had a pretty good understanding of the system, and I found I actually enjoyed talking to people about it. The downside of this was that my activities on the test routines became more and more delayed, and after a while – I can't remember how long – my other activity was formally recognised, and I was officially given the role of "Atlas Support", acting as liaison between the Supervisor team and the sales staff, and with the objective of keeping sales staff and potential customers away from the rest of the Atlas team. I think Dave either took over developing the test routines himself or he decided to do without them, which was probably OK because, knowing the quality of his work and of his mind, there were very unlikely to have been many errors or holes in the central routines.

The only other activity I remember clearly being involved with at that time was the design of an Atlas programming language. This was not a high level language like Fortran or Algol, which came later, but a more user-friendly kind of assembly language, looking ahead to the time when users would be getting on to the machine. It was called Atlas Basic Language, or ABL. I could not have told you when I did this work, except that by chance I recently came across a link via Google to a document entitled "Ferranti Atlas Basic Language (ABL), Provisional Description, by B M M Hardisty, dated 1st June 1962", filed in an organisation called the Centre for Computing History, which I had not previously heard of. So I must have done this work little more than six months after joining Ferranti's. Weird! The actual programming for ABL was done by Jim Doughty, with whom I had worked on the specification. And so far as I know ABL was used for several years on both the London University and Harwell Atlases. I also found a reference to it in the Harwell archives.

By the autumn of 1962 I must have been working more or less full-time on my role in Atlas support, because I don't think I was doing any more active programming for the Supervisor at that time. It was because of this role that I was one of those allocated to attend the inauguration of the Atlas Computer in Manchester in the December, which is of course the occasion we are commemorating this week. This was to be a very high profile event. For those of us coming from London (staff, customers, public figures etc.) this was to start with a grand dinner in a London hotel. I think it was the Hotel Russell, I can't be sure, but it was certainly in Russell Square.

The plan was that, after the dinner, coaches would take all the staff and guests to Euston Station, where a special over-night sleeper train was to be laid on to take us to Manchester for the official inauguration ceremonies the next day. You must remember that, at that time, this was still the old Euston Station, and that it was before electrification. The fastest trains to Manchester took around four hours, and many of them took much longer. So to have an over-night train, with sleeping cars, was not that unreasonable. The plan was that the train would travel through the night and arrive at a siding not far from Manchester in the small hours, where a dining car would be hitched on, filled with the stuff and staff to serve the guests with breakfast before proceeding to Ferranti's at Wythenshawe for a tour round their factory there during the morning, followed by lunch at a hotel in central Manchester, and then the main ceremony at the University.

It didn't happen quite like that.

Yes, we had the dinner, and yes, the coaches took us to Euston Station. But when we got there we were told that the train was not there for us, the engine that was supposed to pull it had broken down, and they were getting a replacement. In the meantime we were invited by the Stationmaster to wait in his office. Because of the occasion, he was all dressed up – like the Fat Controller. But of course his office was quite a modest affair, and certainly not big enough to accommodate in style the large number of people he invited. So we sat on the edge of his desk or stood around until eventually (I cannot remember how long, but it may well have been past midnight by then) we were told that the train had been delivered and we could board it. I went straight to the sleeping car and to bed, but I learned later that many others went to the bar coach for drinks, except that apparently there were initially no lights in that coach and no running water, and the staff had to rig something up from an adjacent carriage. So it was not exactly a five star occasion. Eventually I suppose everyone went to bed. I remember at one

point in the night noticing that the train was stationary, and that at another it was going in the reverse direction! What I do know is that long after the time when we should have arrived at the siding to pick up our breakfast coach, we were still travelling at speed. I learned later that what had happened was that the relief engine had also broken down, that a breakdown engine had been despatched to tow us off the main line and on to a siding to await a second replacement engine, and that was why we had been stationary and going in the wrong direction. Anyway eventually we got to the siding where the breakfast car was waiting. It was hitched up, and in due course we all had breakfast. But by then it was so late that the morning tour was ditched and we went straight to the hotel in central Manchester for the lunch. And, as far as I can recall, the rest of the day went off without any more serious hitches.

My work in Atlas Support must have been expanding, because at some point in 1963 I started taking on staff to work with me. I think the first person I recruited was John Deas, who had taken over the special purpose computer in the Engineering Labs at Cambridge that I had worked on, and who had designed and overseen its replacement with much more up-to-date and reliable electronics. Anyway all this was overshadowed by the take-over of the Ferranti Computer Department by ICT in the autumn of 1963, and the reorganisation that followed. Chris Wilson, who had been in charge of Ferranti computer sales activity, became manager of what was called Sales Division B in ICT, and which was essentially the part of ICT which sold computers – Division A being largely the former ICT sales people who were still largely selling punched card equipment. Under Chris there were three so-called sales regions, which did not cover regions in a geographical sense, but which were called functional regions. There was a Region, whose name I forget, which looked after major industrial companies, there was Government Region, which looked at central and local government's computing needs, and there was what was called Universities and National Research Region, whose name is more or less self-explanatory. The thinking behind this structure was, I think, predominantly so that one so-called Region would look after all the immediate Atlas customers and prospects, allowing all the knowledge, experience and expertise on Atlas to be focussed there. This Region was run by Dave White who had joined Ferranti's earlier in the year. Under him there were two so-called sales areas, one looking after universities under Henry Goldberg, and one under Pat Ronaldson, called National Research Area. I made up the management team with the title of Support Manager, with the responsibility for looking after the technical needs and interests of all customers and potential customers across the region. And, as it turned out, with responsibility for looking after the interests of potential Atlas customers anywhere else across

the company. And so it was that I got involved with all kinds of potential Atlas customers across the world.

Another part of this reorganisation meant that the Atlas Supervisor team was transferred to a Programming Division under Peter Hunt. So I was severed from direct involvement with software development from that time, and I don't think I ever wrote another serious piece of software in my life!

I did innumerable presentations about Atlas, and went alongside sales staff on innumerable visits to potential customers, as well as managing the needs of the existing customers, which at that time were just London University and the Harwell, so-called Chilton, Atlas. And of course in 1964 the ICT 1900 Series was launched, and I took on responsibility for supporting that too.

Amongst my international trips was one to Pittsburgh, supporting a sales pitch to Westinghouse, who were at that time looking for a major upgrade to their systems. We concluded that they needed something like three Atlases and five large 1900s to meet their needs, and we found that this project alone would have taken a whole year's output from the West Gorton factory. I learned from the Bernard Swann report that apparently Westinghouse were very close to placing an order with us. My goodness that really would have put us on our mettle and would have transformed the company or killed it.

Another visit I made was in the UK to visit the person in charge of computing at Rolls Royce aero-engines. This was on an enormous industrial estate in, I think, Derby, where they made aero engines for aircraft manufacturers throughout the world. On this occasion we were accompanied by the local regional sales manager from Division A, his salesman, and also by Basil de Ferranti. We had lunch in a local hostelry, and at one point Mr de Ferranti urged the Rolls Royce man to consider supporting the British computer industry by buying British, i.e. from us. The Rolls Royce man became very angry at this and said to Mr de Ferranti that they sold aero engines to the world not because they were British but because they were the best, and he urged ICT, as it was by then, to do the same. After Mr de Ferranti had left, the Rolls Royce man turned to the local sales manager and said to him words to the effect of: "I never want to see that joker again."

One of our major pre-occupations at that time was estimating or measuring the speed of Atlas for various applications. This was clearly very important for all the customers, since they

were usually looking for the fastest machine for their money to execute their particular programs. When they were putting a project out to tender, they often supplied a program and asked the various manufacturers, who were being asked to bid, to run these programs and time them. Of course it was rarely quite that simple, because the exact conditions under which the programs were run could make a vast difference to their speed, and so there was often a lot of negotiation around the running. Most often there was a central part of the program which dominated the execution time and these loops were given to manufacturers with very precise instructions as to how they should be run.

One such example with which I was involved concerned AWRE Aldermaston. They had historically had wall-to-wall IBM computers, partly because they were working closely with their American counterparts, and no doubt used a lot of the same programs. However on this occasion we were in competition with IBM with their 7030 computer known as STRETCH. They were of course under pressure from HMG to Buy British, as were all public sector bodies, and with Atlas we were in with a real chance. Their chief programmer, one Alick Glennie, personally supervised the running of their loop on Atlas, and also ran it himself on STRETCH. When his analysis was complete, he came back with his conclusions, which he put very succinctly, using a motor racing analogy:

“STRETCH is faster on the straights but Atlas is faster on the corners.”

This was very interesting and also very instructive. You see in many ways STRETCH and Atlas has very similar designs. They both arranged to execute the various stages of several instructions simultaneously in order to achieve maximum overall speed, although STRETCH actually worked on even more instructions at a time than Atlas. However, when Atlas came to a branch instruction it did not access or begin to execute any further instructions until it had worked out whether the branching was to be successful or not. STRETCH on the other hand assumed that the next instructions in sequence would be executed, and if the branch instruction was successful, it then had to unscramble and clear out the partially executed instructions it had begun. This was particularly important on the innermost loops of many programs, where the branching was most usually successful until some particular condition had been satisfied. But of course Atlas would lose out if most branches failed most of the times. This would more often apply to tests in outer loops, where, because they were encountered less frequently, they would be less significant. So I think we more or less tied on points. In the end they bought

both a STRETCH and, subsequently, also an Atlas 2. So I guess honours ended up more or less even on that one.

Talking about Atlas 2, I had much less to do with it than with Atlas 1. It was a collaboration largely between Cambridge and Manchester, and there was a separate Supervisor team in ICT based at Lily Hill House in Bracknell under Chris Spooner, who sadly died before the work was complete. I had one of my team closely involved, but by then I had much broader responsibilities, and did not personally get so closely involved.

So what are my overall impressions of that stage in my career? Firstly that the Atlas design was brilliant, and brilliantly executed. Secondly that we were a very small fish in a very large pond, and we didn't have the resources or the economies of scale to compete in the big bad world, despite the product's brilliance. Not enough of the company's resources were, or indeed could responsibly have been, devoted to such a large machine, for which there was such a relatively limited world-wide market. It would have needed enormous investment over a long period of time to ramp up the scale of its activities, with, of course, no guarantee of a good return. So objectively it was right for Ferranti's computer department to merge with ICT, and for ICT to concentrate on more mainstream computers. But it was sad that, in all of this, Atlas got side-lined, and in effect the company dropped out of the big machine race for a while. It tried to get back in later, but by then we were overshadowed by much more powerful systems.

You may be asking what I actually contributed to Atlas. I have been asking myself that question, and I have come to the conclusion that the answer is: not very much! I joined the team too late to make any contribution to its design, even if I had had that capability, and in the end I contributed almost zilch in the way of software. And moreover almost none of the sales campaigns to which I contributed actually resulted in a sale! Of course I helped in getting the first systems off the ground, attending project meetings, following up actions etc., but nothing very ground-breaking there; just very hard, relentless work. I suppose my main achievement, if it is one, is that I relieved Dave and the rest of his team of the burden and distraction of spending oodles of time with potential customers and others, and thus enable the project to be delivered faster than it might otherwise.

But, for me, the main result was that I became a kind of Mr Atlas for the sales organisation. Any queries about Atlas from a technical point of view tended to be passed to me, and if I couldn't deal with them, I passed them through to the technical staff. This in due course meant that I became thought of as the "expert" in the sales organisation on large systems. I

represented the sales organisation in any planning groups concerned with future large systems, whether operating systems, or complete systems. And this led in turn to my being transferred to central planning in Putney to look after large systems, and eventually to join the team that planned the 2900 Series. That was an amazing experience, because almost everyone who had ever been involved in designing any computer in the UK was there. And during that time I was also involved with negotiations regarding potential partnerships with CDC in the US, Siemens and Nixdorf in Germany, CII in France, and Philips in the Netherlands, all of which came to nothing in the end.

But that's another story.

So overall I had a fantastic time in Ferranti/ICT/ICL. And it all started because Dave Howarth decided he didn't want to spend so much time talking to potential customers. So thank you Dave for that.

I'd like to finish with a short anecdote. I was not personally involved with it, but I heard about it shortly afterwards. It concerns one Conway Berners-Lee, the father of Sir Tim, the creator of the Worldwide Web. Conway is, as far as I know, still alive and well at the age of 91. He was employed by Ferranti/ICT/ICL and was an early user of Atlas and a good friend and colleague. He was also notorious for being very absent-minded. He had apparently been up to Manchester on one occasion to run some programs on Atlas. When he had finished he left with several colleagues to get a train back to London from Piccadilly Station. When they got to the ticket barrier, he could not find his return ticket and had to buy a new one. That was fairly typical Conway, but when he got home his wife's first words were: "But where is the car?!"

Brian Hardisty, 5 December 2012