


50th Anniversary of the Ferranti Atlas computer, December 2012.

Atlas archival interview number 1: Supervisor software.


Transcript of a 40-minute audio interview with D J Howarth (DH) and M T Wyld (MW), conducted in Manchester by Simon Lavington (SL) on 6th December 2012. David Howarth and Mike Wyld were intimately involved in the design and implementation of the Atlas Supervisor (*the Operating System*).



(left-to-right): David Howarth, Mike Wyld, Dai Edwards and Yao Chen on 6th December 2012.



David Howarth in 1962



Mike Wyld in 1962

Biographical notes for David Howarth and Mike Wyld are given at the end of the transcript.

Finally, at the end of this document is a post-script containing further details of Supervisor restart points, based on information given by Mike Wyld in early January 2013.

Transcript.

SL: This is the morning of Thursday 6th December 2012. I'm Simon Lavington and I'm interviewing David Howarth and Mike Wyld. First of all, Dave: 'hello, Dave'

DH: Hello, Simon.

SL: That's a good introduction, and 'hello, Mike'.

MW: Hello, Simon.

SL: Excellent. Well, what I'm going to do, I hope, is mostly to be in the background and prompt with suggested questions for conversation. I wonder whether we may start back in 1960, '61, '62. Atlas was showing signs of coming into life, although of course the final inauguration was not until December 1962. In the run-up to that, there was frantic work on the Supervisor. At the same time you must have been aware that the competition, IBM, was working on their STRETCH computer and the 7090 had come into being. So my first question: When you look back at the facilities that the Atlas Supervisor had, and the facilities offered by IBM on their STRETCH as an Operating System, how did they compare?

MW: Well, let me start. I guess it was around mid, er, early 1961, I was, um, David and Stan Gill who I was working for, said they wanted to look at a Fortran compiler. So I was actually sent off to IBM – I guess it would be about Easter, Easter '61 – and I was there for two or three weeks on a Fortran course and in general sort of being a friendly spy, I suppose, to see what was going on. And the impression I got that, er, was that STRETCH I don't think we had any access to, but the feeling was that it was going to be really not much more than [the Ferranti] Mercury. It was more than a single-user [system] but it was very much, er, a sort of one or two users and they'd split the time in some way or scheme according to some formula or other. But it was really just a case of switching to that program whilst this one was doing a few disk transfers, er, drum transfers, er was it tape transfers it was in those days. But you got the impression that it was all very stylised and not tremendously er, you know, it was really overlapping programs rather than anything else. That was the way – I came away with the impression, er, ...

SL: [at 03:11elapsed time] Dave, can you add anything to that? What did you feel?

DH: Er, no, I get very much the same impression that, whereas Atlas was a genuine multi-user system, the set-up on STRETCH was very much an outgrowth of the old Fortran Monitor System

MW: Yea....

DH: ... which used to run on the 7090s and so on, er, which, I mean, was a step ahead of one user at a time ...

MW: Exactly, that's right, yes, ...

DH: ... but nothing very much elaborate about the multi-user set-up. It was very much an outgrowth of the old Fortran Monitor System that had existed for years on the 709, and the 7090 and then the 7094.

MW: I think we talked about it at the time as an overlapping system rather than a multi-user system ...

DH: I think that's fair comment.

SL: [04:19] So, in embarking on the Supervisor, I guess you felt that really this was new ground.

MW: Very much so, oh definitely ...

DH: As far as I'm concerned, very much so.

MW: It was going from a sort of push-bike to a, you know, a Rolls Royce.

SL: [04:37] So there was no point in talking to anybody else. You were in the lead: is that the sort of impression?

MW: I don't think, it was not a case of no point in talking to anyone else but it was a case of you just felt it was a quantum leap, really.

DH: Yes, indeed; I think that's fair comment.

MW: And also there was a sort of secrecy about it, because, er, certainly Tom I got the impression, and Ferranti, did not really want people to know what was going on. Tom in particular, I think, with some of the patents and stuff like that, er, everybody seemed to be producing patents about once a week for everything and it was very much a sort of ... And of course that goes back to the F C Williams thing, Freddie Williams, patent ... We used to joke about his coffee machine and his tea maker and things like that ...

SL: ... and egg boiler ...

MW: ... his egg boiler, indeed, he'd sort of patent on a Tuesday afternoon for a bit of fun. And er... So patenting, and copyright, things like that, certainly were in the Ferranti mind and certainly I think was in the University mind. Yes.

SL: [05:45] Thank you. May I switch to a sort of administrative topic for a minute? I believe you both reported formally to, well, first of all Stan Gill mainly and then to Hugh Devonald as your line managers. Did either of those two gentlemen, either Stan Gill or Hugh Devonald, take much active role in code-cutting, in actually what happened

DH: Well, to a very large extent as far as I was concerned they were kind enough, if you like, to leave it to me as far as the details were concerned. I would report to them whenever anything significant happened, or say once a week or something of that order. Their line management arrangement, as far as I recollect, was very good and very fair. They provided me with the resources that I needed, at the time when I needed them, and so on, ...

MW: Yes, that's my recollection. Stan Gill interviewed me for the job. Then Hugh in fact, who was a good friend, but he was very much administrative manager: pay, rations and, er, you know. I think he did my appraisal a couple of times but he was, you know, 'How are you doing?' Tick, tick, tick, and move on, type of thing. Stan was, I guess, more the technical man but he just seemed to be involved in everything so, again, I think, his contribution as far as I was concerned ... I sat outside his office and er he'd sort of give me to er the coffee machine and 'How's it going?' sort of thing. Hugh Devonald got involved in the Orion project. He seemed to be doing things with that. But I don't think that they wrote any code, or anything like that.

DH; No, I know. Sure, sure.

MW: Stan was very much a sort of politician, I guess.

SL: [08:04] Pursuing this line of colleagues and who you reported to: when you both got involved with the Supervisor I think the man already in post, as it were, was Bruce Payne, who was an academic at Manchester. If I understand it correctly, when you arrived Bruce was the Supervisor person on Tom's team....

DH: That's correct.

MW: Yes, absolutely. ...

SL: ... and over the years, the next one or two years, it seems to me (from the outside) that Bruce rather took ... stepped back a bit. Can you comment on how you found Bruce? The working relationship, that sort of thing?

DH: Oh, yes. As you say ... When I was recruited, as it were, there was some question in Stanley Gill's mind as to whether I should be assigned to the Compiler Compiler project, which was running in parallel with Tony Brooker, or whether I should add my two pennyworth, as it were, to the Supervisor work. And at that time I think the ICL people, the Ferranti people, didn't have anybody full-time working on the Supervisor side of things – they left it to Bruce Payne. So, when I was first recruited, it was as a co-worker with Bruce Payne: he reflecting the University approach and myself heading the team representing the commercial approach, if you like, of the Atlas development. So I worked with Bruce Payne, oh for, as you mentioned, a number of years. He was a close personal friend, apart from anything else. He was a good sound man.

MW: You shared an office with him for a time, if I remember rightly ...

DH: Oh, yes, I did.

MW: We came up [to Manchester] and were very short of space. And I remember David and Bruce sort of – I can't remember whether you even had a desk each or [sat] opposite sides of a desk ...

DH: That's right ...

MW: ... I can't remember now but certainly in the same room ...

SL: The impression I got was that, as the project got more and more intense, Bruce was quite content to let other people cut the code, as it were ..

MW: Yes, possibly ... I don't remember anybody else being involved really, from the University side, in any ...

DH: No, I ...

MW: ... somebody worked on the peripheral software ...but I can't think that there was anything really involved in the detail -- scheduler and Supervisor and so on ..

SL: [11:16] Now, although you didn't formally report to Tom, Tom Kilburn, I guess his presence was overpowering in the sense that he was ... would it be right in saying, his was the main driving force of the whole project?

DH: I think that's fair comment.

MW: Absolutely, yes.

DH: I think that's fair comment. One was always conscious of the fact that Kilburn was in charge. He was the one who knew everything that was going on.

MW: He knew everything ... what you were doing ... why are you doing that? Or that sort of thing. Totally, I mean. In fact, I don't know where he actually did any work but he sat in the computer room, is my memory. Almost every time you went in, Tom would be in the corner with somebody or other – with David, Dai, or whatever. But I'm sure he had an office somewhere - I've no idea where it was!

DH: Sure.

MW: He was very much on top of things ...

DH: He was always there, and you could rely on him being there, if reliance was useful, or you could ignore him being there. It was really quite encouraging to know that a man of

his intellectual prowess, as it were, was sat there in the room puffing away at his pipe, keeping an eye on us hard workers, as it were.

MW: And I think everybody deferred to him. I mean, Frank Sumner, you know, what I saw was the senior university software guy - Derrick Morris ... In fact I shared a room with ... Frank and Derrick had got a room and I went in as their 'box in the corner' sort of thing for about three years.

DH: Did you?

MW: So you were with Bruce on one side of the corridor and Tony Brooker was next and I was on the other side with Frank and Derrick Morris ... a little huddle ...

SL: [13:22] I referred [before the interview] to what you might call *software aids*, *program development aids* - I mean tools as we now call them. I guess that in those days there weren't very many tools at all and perhaps just taking a dump of the Supervisor Working Store was the nearest you got to a tool? Tell me, first of all, starting from basics, before you had a lineprinter, if you took a dump of the working memory, which was 1K words, it now had to come out on the teleprinter, didn't it?

DH: Correct...

SL: As octal, or as what?

DH: Probably as octal.

MW: In octal, yes.

SL: But how much paper was that?

MW: Oh, a thousand words ...

DH: A thousand words ...

MW: And you always got all ... My memory was that you couldn't get ten ...

DH: What's your guess, Mike, how long ...

MW: I guess it was about ten pages ..

DH: I would support that. At least that ...

SL: And took some time ...

DH: Yes, of course.

MW: Well, I mean, in a timing situation we were ... I was up here [in Manchester] I suppose most of the time from mid-61 and we used to get machine time in the evening. So, we'd be here all day – we stayed in some dreadful hotel across the road. I think it was about £1 a night, or so. But we stayed there and we used to work all day, literally from nine to six, and then we'd go and have a meal, and then we'd come back at sort of 7 o'clock, and that's when we'd get the machine. So we'd work on the machine from 7:00 to sort of midnight and then we had to go back to the hotel because they locked the door at about eleven o'clock and we couldn't get in if we were after eleven o'clock. So we'd get back to the hotel, but we'd leave ... some sort of night watchman or something .. you'd leave you stuff printing and come and pick it up in the morning. That's my memory of how we worked. I don't know whether that's your ...

DH: Yes...

MW: I mean, we'd go through the night ... We always got Thursday night. Thursday night was our night on the machine for the Supervisor ...

SL: And then back to London by train on the Friday? Or?

MW: David and I ...

DH: Yes, certainly as far as I was concerned my living home, as it were, was in the London area, and so I would travel back, as you say, on the Friday or Friday evening and come up again on the Monday, ready for the fray once again.

SL: [16:11] I've heard it said that - correct me if I'm wrong - sometimes you were so involved that you didn't bother to check into your hotel. Is that right?

DH: [laughs]

MW: He didn't even bother to book one! [Laughter].

DH: Ah well, I mean this is just a reflection of the fact that there was just a lot of work to be done and it was being done at a distance, as it were, and it just wasn't appropriate to set up a room and contacts and so on in the Manchester area.

SL: And this was before, of course, e-mail or mobile 'phone ...

DH: Oh, yes!

MW: And the train was about a five-hour journey. In fact, I remember I worked out the quickest way to get back to London ... I used to go from here [ie Manchester] to a station called Miller's Dale in Derbyshire, which was like a junction, get the train from there to Derby, I think, and then a fast train from Derby to London. That was quicker than going from Manchester to, er, at the time of night when you wanted to go ...and I think I got four trains ...

SL: Steam trains ...

DH: Oh indeed yes.

MW: Oh, yes yes!

DH: Yes. It would have transformed life in many ways in the project if we had had things like mobile 'phones and e-mail addresses and things of this sort.

MW: Oh, and travel by road ... I mean ... We got married in August 63 and in fact came to live in Manchester. My family lived in Birmingham. And I used to leave ... My wife was teaching and I used to leave the school in 3:00 in the afternoon to get to Birmingham, not London, pick her up at 3:00 in the afternoon and if we got home by ten o'clock that was a bonus. Seven hours [driving] to Birmingham. But I mean, you didn't drive anywhere, really.

SL: [18:15] Pursuing this business of no e-mail, no mobile: as Atlas came into full production on the three main sites, the Chilton site, Manchester and the London site, keeping the software, the Supervisor, consistent and up-to-date on all three sites must have caused some headaches? So: talk us through this business of trying to be the guiding figure of three sites.

DH: Well, I was ... this is ... this turned out to be quite a problem, really. I mean, if we'd have been given e-mail, mobile 'phones and so on, it would have been very easy to construct to the same thing: one going at Chilton, one going at London, and the prototype, if you like, running at Manchester. In practice, without things like mobile 'phones and e-mails and so on, in practice there was really quite a lot of work to be done, to ensure that all three stayed roughly speaking along the same lines and at the same level of development. And, er, we managed.

MW: My memory was: I was here in Manchester and I think you, in the end, spent a lot of your time with the London machine, if I remember rightly.

DH: Yes, I did.

MW: And then was it Peter Warne and Mike Bayliss were at Aldermaston – er, Harwell.

DH: Yes, that's right. That's right.

MW: And around that time the Atomic Energy Authority set up an aeroplane [see footnote at end] – it was an old Dakota – and this used to start at whatever the place that was called [Sellafield?], in the Lake District, there was a nuclear station there and it used to come down here and land at Manchester Airport to pick up people from the place – there was an Atomic Energy Authority place at, er, near Knutsford, out that way [Risley?], and then fly on down to Aldermaston and then fly on down to Bournemouth to look after them [at AEA Winfrith?]. So I used to hitch a lift on that 'plane, my memory was, every other

Thursday or every other Wednesday or something. And so I did that, it was quite funny, I did that for about a year, I guess. I said, 'How do you know which 'plane to catch?' And so they said 'You go to Manchester Airport and there will be a lady wearing a large hat with flowers all over it. And you go up to her and she'll tell you where to go'. And I turned up at Manchester Airport at 7:30 on the morning and there was this, well, a large lady with a large hat as well, and she would say 'Would you like coffee or tea?' And so I'd say 'coffee' and then you'd get on ... And this Dakota would arrive from wherever it had started, and you'd get on with a paper bag full of cups of coffee or tea – (there were cardboard cups just coming in). And then you'd get on to the Dakota and fly off down to, er, and then it came back at about 5:00 in the afternoon, and it did that every day. But it only went that way it didn't come the other way so it wasn't any good to them coming up from London or Aldermaston. So we tended ... we seemed to meet ... my memory was that we met quite often at Aldermaston – at er Harwell.

DH: Yes.

MW: But then I think we gradually got to the stage of ... we would come down to London.

DH: Yes, that's right.

MW: We tended to alternate between the three sites. London always had the prettiest operators ...

SL: [22:20] I know we shouldn't do, but let me pursue that theme for a moment. Most of the people involved in the design and implementation of Atlas were male – they were men. There must have been surely some exceptions. I do remember hearing from a lady called Ann Moffatt. Ann worked for Kodak. Kodak had an urgent problem, written in Mercury Autocode, so that in the summer of 1962 Ann was one of a very few users who were allowed time on the Manchester Atlas. She describes coming up to Manchester, staying in the Midland [Hotel], working nights, as the only female amongst a group of men, and how the males were very gentlemanly in allowing her to sleep in the only chair in the Atlas room whilst they slept on the floor. Have you any memories of any ...

MW: I do remember Kodak being an early user and I do have a memory of a lady but I couldn't have told you her name was Ann Moffatt. And I don't remember her from being from Kodak, to be quite honest, but there was a lady who was involved with one of the users. I've now read your note and I guess that that was ...

SL: But that was a rarity, I guess, wasn't it?

MW: Well, the only other lady who was around was Miss Popplewell, Cicely Popplewell.

DH: Oh, Cicely Popplewell ...

MW: ... That's right....

SL: ...who looked after the University's Computing Service ...

DH: That's right....

SL: ... which continued on Mercury, I guess.

MW: There was one super incident where ... There was one of the users had got a program ... and it was somebody working downstairs with Freddie Williams, and this thing [program] kept giving variable results. It was Mac [Graham Maclean] who had written a program and it was, anyway, there was something weird. You'd run the same program ten minutes after and you'd get different results. And Cicely Popplewell came and, she was brilliant on the administrative side, very good to us all, and she was the world's expert on a Brunsviga whirly machine [a mechanical desk calculating machine] and she actually did out these equations on her Brunsviga, which took about three days, by hand, and she proved that it was the equation that was unstable.

SL: Ah!

MW: It was an unstable equation and in fact there was a problem with the floating-point [unit on Atlas] on about the ninth digit or something and this was causing the issue. But it was Cicely on her Brunsviga who proved that Atlas was right! I don't know when it was, but er ...

SL: [25:33] We got side-tracked in a way, because I started talking about software aids and you described the dump of the Working Store. What other aids? I remember yesterday [5th December] in the presentations we had [in the Atlas Anniversary Symposium] the Cambridge man [Barry Landy] talked about TSDS, the Titan Supervisor Diagnostic System. I don't know what that was. Did you have anything that could be called a *Diagnostic System*?

DH: Well, not really. I mean we more or less worked on the, er, 'developing what you need at that moment', in order to help our life along. We didn't have any sort of theory as to how to do it. But, as one example of this, we set about writing the code of the Supervisor and most of that code was in fact written in *Intermediate Input* as it was called. Very low level.

MW: It was a sort of Assembler type. Before ABL [Atlas Basic Language].

DH: That's' right. I remember talking to the Compiler Compiler people and they said, well, 'Clearly you've got a big project on the go here and why don't you develop a language to make life a bit easier?' instead of writing things like *ADD* or *124*, or whatever it was you wrote down, to make life a bit easier. Now it's a good thing we didn't in fact follow their advice because, if we had have done, then the whole team would have turned out into Compiler Compiler people and we would have ended more concerned with how it was written compared with what was coming out at the far end. And what we wanted to do was to bring out the right thing at the far end. So we sort of put up with the fact that, er, we

didn't have all the nice glossy tools that the average user would demand. We put up with it because we wanted to pick up a result at the end.

MW: The other thing ... we sort of ... the lineprinter was an absolute godsend when it turned up....

DH: Oh, yes.

MW: ... and [we] developed a bit of software which, I think, sat in the Fixed Store, if I remember rightly, but it meant that if there was a fault the operators had to just press one button and we got a dump of all the relevant ...

DH: That's right.

SL: Absolutely right, and if you look at the - Eric Sunderland [a skilled maintenance engineer on the Manchester Atlas] tells me - that if you look at the display [at the Atlas Anniversary Symposium] which is in the Kilburn Building today, the handkeys on the console are set up so that you can do a Working Store dump at the touch of a button.

MW: In fact you used to get a display at the top [row of the console lights] which showed, er, the register that was what it was, type of thing, when the fault had occurred. But that was so ... Quite often I'd get a call at two o'clock in the morning to come in because it was, er, something had happened and we go the little, you know, er what they had to do if there was a fault. It was just: 'hit the button, do the dump' and restart. And nine times out of ten it would kick into action again. And then come in the morning and then analyse ... I mean, some mornings I'd have seven or eight there to sit there and wade through. But we kept ... but certainly I kept all the Manchester dumps and logged and listed and then when you'd found a fault you could go back and tick off and say, 'Well, that's caused that one, and that one ...'. And then it gradually got to be two or three a week, rather than ...

SL: [29:58] But still concentrating on the intense period of 1962, when the work-rate demanded of you was, er, ... how did your family, other folk, your friends, have to put up with this? Any domestic issues?

MW: Well, I wasn't married, but, er, it sort of came with the job right from the beginning. When I was working on Atlas, in Newman Street [Ferranti's London headquarters] on the other side of the room were the Orion people, and Nebula. And they got into all sorts of problems with, er, it was the Prudential [insurance company] was the first Orion customer: they were the one that was about a year late. And, to help out, we would work in Newman Street until five in the afternoon or whatever it was, go and have a meal, and then come back, and we'd pretend to be the Nebula compiler till midnight, or something, because the Prudential were writing all these programs and the Nebula compiler was taking hours - well, (a) it was taking hours to produce a program and then the program it produced would run for hours, doing the simplest thing. And they'd got these huge sort of handshake series that they had to go through between every set of instructions. And you could have a simple instruction like, you know, 'Get next record' or something and it still went through

this God knows how many thousand words of code before it would do the next instruction. And really all it needed was about, you know, four words of code. And so we used to literally go in in the evening and run down the programs, see these instructions, and just take all this ruddy code out (excuse me!) ... and take the code out and just put in what was needed, because whilst Nebula was a fantastic - the idea was you'd just take text and put it in – but people didn't program like that. They did say 'Get next record' and what you wanted to do with it, and so, even before Atlas, I was doing that sort of thing - where you - I probably did three nights a week for six months on that ...

SL: Compared with that, Atlas was much more, sharper ...

DH: Yes, that's right.

MW: Absolutely.

SL: [32:49] But, Dave, you didn't get a chance to make any comment, if you wish to make a comment, about family and friends during these times of intense pressure when you disappeared up to Manchester and didn't sleep for five nights. What was the effect on your family?

DH: Well, it was a fair amount of strain involved. Er, I was married just before I joined Ferrantis and pretty well straight after, therefore, after some sort of honeymoon I suppose, I would be vanishing for the week up to Manchester whilst my wife managed in our little flat in London. Well, it wasn't the ideal way of developing family life by any manner of means. I remember my wife was extremely pleased when the Atlas, when the London Atlas came on line because at least I could work a good deal of the week around the London area and see something of my wife and join her for lunch and things – civilised things like that. She was very pleased when we were able to get together, rather than being separated for long periods. So it did have an effect on our family life. I was incredibly grateful to my wife for putting up with it, to be honest, because she had every reason to object to the idea of a husband being absent for week after week and this wasn't her idea of married life, as it were. But we survived that period and, once the London Atlas began to come on line, things were transformed completely. I could work from nine to five, or whatever, and take a tube [underground] to get home. I think it was much more like a normal home life. It's at that time of life that my daughter appeared. She has brought me up here [to Manchester] today.

MW: Well, when we got married in August 63, we discovered on our wedding day that we'd never seen each other on a Tuesday, Wednesday or Thursday! The Tuesday after we were married was the first Tuesday we'd ever ...

SL: It's been delightful to hear all these anecdotes. Are there any concluding comments that you'd like to put on record? Things that you feel perhaps aren't sufficiently recognised? One of you may wish to touch on this business of publication and promotion and things like that. Any comments? Any injustices that you feel need to be cleared up?

MW: Well, I don't feel there are any injustices. I mean, we could ... Basil Ferranti's comment that, you know, 'You'll be very proud in your life to have worked on Atlas'. And only a few years ago, well, I was IT Director of the [National] Lottery and I took on a programmer, a Program Manager, and he'd been in the job about three weeks and he appeared in my office and said 'Are you the Mike Wyld who worked on the Atlas Supervisor?' And his Dad was one of our engineers, as it happened, but it just came up through your working life that it, er ...

SL: You were glad you did it?

MW: Oh, fantastic! Yes, and I think the thing that really, sort of ...the ICL take-over, well the ICT take-over then the ICL, that really knocked the stuffing out of Ferranti I think. And that was really when I think an awful lot of people left in 66, 67, 68.

SL: Dave: anything to add to that? Are you pleased you worked on Atlas?

DH: Oh, yes, yes, yes! Enormously pleased. I just wish I'd been able to spare time to do other things on the side, as it were. It was a full-time job for a time, at any rate, and I'm enormously pleased to have worked on it. I don't know what life would have been if I'd sat back and accepted some more menial job in Ferrantis, for example. I'm very proud, in fact, of what we achieved and enormously pleased to think that I was involved in that achievement.

SL: There is a book, I think it's called *Significant Advances in Operating Systems* by Per Brynch Hanson I think published a few years ago [actually 2001] and they do make credit to the Atlas Supervisor as one of the outstanding achievements in the world of Operating Systems...

DH: Yes, that's right.

SL: ... and I think that those who know do recognise ...

MW: And I think that so much of it goes to David, because he was the guy who really, I mean ... there was some very brief papers at the beginning that really pointed the direction and the thing was steered through and, as you say, Tom led the project but David was the man who really held the Supervisor together and steered it through.

SL: So, if I may say so, when I was talking with Tom when [in 1974] I was writing that little booklet about Manchester computers, Tom was most emphatic that you should be mentioned, Dave.

DH: Really?

SL: Tom definitely knew who to credit all the Supervisor ...

MW: Absolutely.

DH: Really ...

MW: It's wonderful ...

SL: Well, thank you so much, both of you, it's been delightful to hear you talk. Ah, Beth ...

[A few minutes before this, Beth from the Manchester Student Radio had come into the studio as our allocated session was coming to its end]

SL: ... Beth, any questions, anything would you like to ask?

Beth: You've covered everything, I think. I was really interested in your family life ... lovely listening to you two together ...Lovely to meet you ...

DH: Good.

SL: Well, thank you very much indeed. It's been absolutely fabulous.

MW: Thank you.

DH: Thank you.

Footnote to page 8. F R A Hopgood, who worked at Harwell during the relevant time, has said: "The Dakota used to land at Abingdon airfield on the way down from Manchester to Bournemouth (it linked Risley, Harwell and Winfrith). It was an early UK network transporting card decks to the 7090 at Risley from Harwell and Winfrith and lineprinter output on the way back".

Elapsed time at end of recording: 39:55.

MP3 file size = 91.4Mbytes. Recorded on 6th December 2012 at the Fuse FM Radio studio in the Students' Union, University of Manchester, M13 9PL.

Biographical notes.

D J Howarth.

Born

London, 29th March 1931.

Education

B.Sc. Physics, Imperial College, University of London, 1951.

Ph.D. Imperial College, 1953.

Employment

Massachusetts Institute of Technology: 1953 – 1954.

SSO at the Royal Radar Establishment, Malvern, 1954 – 1960.

Ferranti/ICT/ICL: 1960 – 1970.

Professor of Computer Science, Institute of Computer Science,

University of London, 1970 – 1973.
Professor of Computer Science, Department of Computing,
Imperial College, University of London: 1973 until retirement in
1996.

Summary of involvement with computers, especially with Atlas.

Whilst a Physics research student at Imperial College, David used the Ferranti Mark I at Manchester to handle band structure calculations in solid state physics. He then spent a year at MIT using the Whirlwind system for similar calculations, followed by a period at the Royal Radar Establishment (Malvern), using the TREAC machine for problems in mathematical physics.

Upon joining Ferranti Ltd. in 1960, David married and moved to London. He was based at Ferranti's Newman Street offices in London where he became the first Ferranti programmer to work on the Atlas Supervisor. Quite soon David became the leader of the Supervisor team, which numbered about six people. He remained the principal Atlas Supervisor expert well into the late-1960s, by which time he had also become involved with the design of ICL's *New Range*, the ICL 2900 series of computers.

In 1970 David moved from industry into academia where, amongst other things, he lectured on Operating Systems.

M T Wyld.

Born	Coventry, 28 th April 1938.
Education	B.Sc. Mathematics, University of Manchester, 1959. Diploma for Advanced Studies (Computer Science), University of Manchester, 1960.
Employment	Ferranti Ltd. (later ICT), 1960 – 1966. CPA (Later English Calico, later Tootal), 1966 – 1976. ICL, 1976 – 1993. Camelot, as an ICL-nominated Director, 1993 – retirement in 2002.

Summary of involvement with computers, especially with Atlas.

Mike Wyld first learned programming (Mercury machine code) in 1958 whilst a student. At Ferranti, Mike started off in Ferranti's Newman Street offices in London, where he gave Mercury programming courses and wrote applications programs for customers. In early 1961 Mike began to work with Dave Howarth on Atlas software. His initial tasks included the writing of a series of Atlas simulators running on a Ferranti Sirius, to estimate critical times of interrupt routines, etc. From about mid-1961 Ferranti set up two teams on the third floor at Newman Street. On one side of the room was the Orion team and the other was the Atlas team. David Howarth pulled together an Atlas Supervisor team of about half a dozen people. Peter Mumford worked on the Magnetic Tape systems, Jim ?? was

writing the ABL compiler, Peter Warn, Peter Jones and Mike Wyld all worked on the central Supervisor. Mike Bayliss and Brian Hardisty were working on documentation and interfacing with rest of Ferranti and customers. Bruce Payne was the main contact at Manchester University and various members of the team travelled up or down on a regular basis. "This was the first time that we all became dedicated to Atlas as our prime job".

In early 1962 it became clear that Ferranti's London team needed to spend more time in Manchester. Peter Jones and Mike became the main travellers, spending weeks about between Manchester and London. By Autumn 1962 Peter & Mike were spending 90% of their time in Manchester, usually working a couple of night shifts per week as it was the only time that they could get test time. Mike says that "David Howarth came up most weeks and didn't even both booking a hotel room – he just worked all night and then slept on the train back to London".

At the end of 1962 Peter Jones and Mike moved full time to Manchester. Mike shared a small room with Frank Sumner and Derrick Morris on the top floor of the Electrical Engineering Department's Dover Street building, and worked closely with Frank Sumner, Derrick Morris and David Howarth on the Supervisor development. Mike was the resident system software specialist in Manchester from mid-1963 until 1966 when he left Ferranti (or ICL as it had become). Up to 1966 Mike's job consisted of working with David Howarth on any new Supervisor functionality required and investigating, documenting and correcting all system failures on the Manchester Atlas. This was a full time job as there were several system "incidents" per day. As the other Atlas systems came online, David Howarth looked after Harwell and London systems and Peter Warn was at Aldermaston.

In 1966, after the ICT/ICL takeover, Mike moved to CPA (Later English Calico, later Tootal) based in Manchester to install their first computer, an ICT 1904A. CPA was the 40th biggest company in the UK and Mike held various titles including *Chief Programmer* and finally *System Development Manager*. In 1976 Mike returned to ICL as Sales Support Manager/Project Manager and worked with most of their large Customers. In 1993 Mike was one of two ICL nominated Directors of Camelot, the Operator of the National Lottery. He was Director of various departments until finally being Director of IT until he retired in 2002.

The post-script follows on the next page.

Post-script to the interviews conducted at Manchester on 6th December 2012.

Two challenges encountered when developing the Ferranti Atlas computer.

This note describes two examples of fault-finding issues that posed interesting problems at the time. It draws on information given in the interviews with four Atlas pioneers [see reference 1] and on subsequent correspondence with the interviewees.

The issues are each examples of what could be described as: *Things can change whilst you're not looking*. In 1960 – 1963, when Atlas was being developed, the two issues were relatively novel and were brought sharply into focus by the novel nature of the Atlas high-performance system. In particular, Atlas had:

- no CPU clock – (control was asynchronous);
- many on-line peripheral devices, of various speeds and crisis-times;
- about 150 distinguishable causes of Interrupt;
- a multi-programming, multi-tasking Operating System.

For an introduction to Atlas and a list of references, see:

<http://curation.cs.manchester.ac.uk/atlas/docs/The%20Atlas%20story.pdf>

Issue (1): arbitration between asynchronous events.

At the Atlas hardware level, it was necessary to decide between two or more independent units that were signalling a request for a single resource. Arbitration was made on the basis of fixed priorities. Problems arose when two or more asynchronous requests reached the decision-making circuit very close together in time. This could cause a narrower-than-normal pulse to be produced – for example 50 nsec. instead of 100 nsec. wide. Narrower pulses might or might not convey sufficient energy to change the state of a decision flip-flop (bistable). Under adverse conditions, the flip-flop might even hover in a *metastable* state, only finally reverting to the standard *set* or *reset* states after a certain *settling-time*. This phenomenon was recognised by the Atlas design engineers at Manchester around 1959/60, at the time of developing the Atlas Pilot Model [ref. 2]. Dai Edwards also remembers discussing the issues with David Wheeler of Cambridge University. Dai says that it was not known whether anyone else, world-wide, had yet drawn attention to the *settling-time* problem. [ref. 2].

As is mentioned by Dai Edwards in the interview [ref. 1], one important Atlas area where the problem was known to manifest itself was in arbitrating between three relatively high-priority autonomous units (the central processor, the drum system and the magnetic tape system) when they requested access to the main core store. Some experiments were therefore carried out at Manchester, on the basis of which an allowance for *settling-time* was made for Atlas that appeared to give an acceptable mean time between failures (MTBF) consistent with the known maximum rate of requests to access the main core store. In the limit, of course, there was no known physical circuit that could guarantee unambiguous decisions in a reasonable finite time.

Ten years later, in about 1970, engineers at Manchester re-visited the *settling-time* problem because it had become critically important during the design of the MU5 high-performance computer, the successor to Atlas. More experiments were carried out, resulting in the derivation of formal mathematical relationships between settling-times, the gain-bandwidth product of the decision flip-flop and the consequential likely mean times between failures (MTBF) for a given rate of requests [ref. 3]. For a desired MTBF, the faster (ie better HF performance) the flip-flop, the shorter the required settling-time.

Issue (2): Restart points.

When writing code at the centre of an Operating System, particularly Interrupt level subroutines, one of the most important things to decide is where to locate *Restart Points*. Any subroutine in the Atlas Supervisor could be interrupted by a more important Interrupt subroutine (for example, those associated with the drum system or the magnetic tape system). If this happened, then control was in due course returned to the original subroutine at a *Restart Point*. This meant that some instructions could be unexpectedly carried out twice, and care was needed to ensure that the contents of some index registers did not get corrupted. This was particularly important in the case of single-bit switches (of which there were several). If a Supervisor subroutine was interrupted after setting a switch and before resetting the *Restart Point*, it was possible to find oneself in a subroutine with an “irregular” switch setting. For Atlas, the general rule was always to use Boolean instructions (rather than arithmetic instructions) in switch settings.

In his talk at the Atlas 50th Anniversary Symposium in Manchester on 5th December 2012, David Howarth touched on this point. The general type of incident that David discussed, where a peripheral fired off a loose signal and the peripheral fault-handling software mis-handled it, was difficult to trace. The faults were usually inconsistent and were notoriously difficult to reproduce. Mike Wyld remembers [ref. 4] spending ages looking for a problem caused by a dodgy *Door open/closed* switch on a magnetic tape deck that caused an issue in some unconnected part of the Supervisor system. A less-than-rigorous understanding of the importance of the *Restart Point* technique was the cause of many intermittent incidents in the early days of developing the Atlas Supervisor.

Someone once said that if a problem was consistent it was probably software and if inconsistent it was probably hardware. On Atlas this was no longer true. Variations in the timing or order-of-occurrence of Interrupts could cause Supervisor software to behave differently on different occasions. If an interrupt level subroutine had a very high priority (say dealing with a specific magnetic tape situation) it might very rarely get interrupted, and there were several situations where routines with incorrect *Restart Point* settings ran for many months without causing problems. Whilst the code might have been logically correct, incorrect positioning of the *Restart Point* was the cause of some interesting problems.

This *Restart Point* concept was used throughout the Atlas Supervisor code and was an important aid to fault finding, since the last *Restart Point* setting told you where the software had been when a fault occurred. The first thing one did when called to look at a suspected Supervisor fault was to check the *Restart Point* and the console display was usually set to that address [ref. 4].

References.

1. Two archival audio interviews: (a) with D B G Edwards and E C Y Chen, and (b) with D Howarth and M T Wyld, 6th December 2012.
2. Discussion with D B G Edwards, 9th April 2013.
3. D J Kinniment & J V Woods, *Synchronisation and arbitration circuits in digital computers*. Proc. IEE vol. 123, 1976, pages 961 – 966. See also: D J Kinniment, *Synchronisation and arbitration in digital systems*. Wiley, 2007. ISBN: 978-0-470-51082-7.
4. M T Wyld, e-mail sent to Simon Lavington dated 28th December 2012.

Simon Lavington
lavis@essex.ac.uk

23rd May 2013.