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

Atlas archival interview number 2: Hardware.

Transcript of a 58-minute audio interview with D B G Edwards (DE) and E C Y Chen (YC), conducted in Manchester by Simon Lavington (SL) on 6th December 2012. Dai Edwards and Yao Chen were intimately involved in the design and implementation of Atlas hardware.

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

Dai Edwards *(Centre)* in 1961. Dai is holding an Atlas printed-circuit board.

Yao Chen in 1962. Tom Kilburn is on the right. *(Photo Daily Telegraph, 23rd Feb. 1962)*

Biographical notes for Dai Edwards and Yao Chen are given at the end of the transcript.

Finally, at the end of this document is a post-script containing further details of asynchronous circuits, based on information given by Dai Edwards in early 2013.

Transcript.

SL: [00:26 elapsed time] Well, this is Thursday 6th December 2012 and my name is Simon Lavington and I'm interviewing Yao Chen: Hello, Yao …

YC: Hello, Simon.

SL: … and Dai Edwards: Hello, Dai.

DE: Very pleased to meet you, Simon.

SL: Right. What I'd like to start with is your comments on Tom Kilburn. Of all the people associated with Atlas, in terms of publications, etc., one constantly comes across the name Tom Kilburn – and quite rightly so. But what I don't know, and our listeners will not know, is: What was it like working for Tom? How did you find Tom as, if I may say, a group leader? Yao: would you like to comment on Tom?

YC: I have no words … how to describe how I feel about Tom. And, I mean, he just, to me, was the greatest man ever. And, either as a project leader, as a professor, everything and, er, we worked under him. We never felt, or I never felt, any real pressure he put on me or to any other colleagues. But everybody worked as they, you know, according to their own true feelings – not as a piece of work but just as a way of life, if you like. So, to me, Tom was everything.

SL: Dai, I don't know whether you'd like to comment because you clearly knew Tom for longer, historically.

DE: Yes, well I knew Tom very early on when in fact, whilst he was an experienced engineer from RRE [at that time TRE, the Telecommunications Research Establishment], when I joined the group he was just completing his Ph.D. So my first contact with him was to actually read his Ph.D. But I always found him supportive, for example, er, me being promoted to Assistant Lecturer after one year. I had got my M.Sc. but he was very supportive in getting me that position, so to speak, which I was very lucky to get. Later on, when we'd completed the expansion of the machine and things, there was a proposal to write a paper. And I met up with Tom and two other people and we were each given a section of the paper to write, and Tom said he would put it all together for us…

SL: This was on the University Mark I?

DE: It was, yes. And, er, we went to the meeting. It was all done on time, we all went to the meeting in Tom's office, which was then in the new building that had been with the Mark I in, and he sat down and he said: 'I've read each of your sections. They're so different in style that I would find it very difficult to just combine them simply, and I've just been invited to go to America, and so I haven't got the time to devote to that and so there's only one thing I can do …' And he threw all the contributions into the bin and left the room! Now, one of the other people was Brian Pollard [Ferranti engineer] and the other was Geoff Tootill, both of whom were really senior to me. And I sat there on my own in the room, thinking 'what shall I do?' And of course, in the University in those days it was 'publish or perish'. So I picked all the sections out of the bin, went away, re-wrote the paper, and submitted it to the IEE [Institution of Electrical Engineers] – without any further

permission, but in accordance with what we'd agreed about how the names [of authors] were going to be specified and the title and so on, which had all been settled. But I went ahead and did it on that. And of course Tom was very surprised when he came back from his visit and after a few months had gone by, I think, he got a thing from the IEE saying that the paper had been accepted. [Laughter] So this was my real first experience, you know, of, really, Tom being very direct and very purposeful, also very determined in the actions that he was going to take.

SL: [05:07]. Later on in life, at least certainly when I knew him, he always used to take August off and I think his regular holiday retreat was Blackpool, was it not? It surprised me, at a time when I and other junior lecturers took our statutory two weeks off, er, it rather surprised me that Tom would take the whole of August off. Any comments about that?

DE: Yes, well, you mention Blackpool. That was certainly one of his venues. But he did also go to Butlins. And he recommended that I follow that route, and I did for the odd occasion. And I must say I enjoyed it with the small families, so to speak. But I think Tom's month off in August, er, he did go for a holiday, he did always return in time to get the start of the football season, where was his main interest. But also he came back with a plan to set off for the following year, where he had mapped out really what we were all going to be paying attention to. Now, it wasn't that he said 'This is precisely what you're going to do', it was that there was an indication of the route which he was happy to talk with you about, and get agreement on how to proceed. But he certainly spent quite a bit of time in August planning his work for the next year.

SL: [6:39] Turning to you, Yao, if I may, my endearing memory of Tom and Atlas is Tom sitting there, with his pipe, not saying much but being a constant presence. How … Tell us a bit about this.

YC: No, I think it was very inspiring because he knew what was going on. He never asked many questions, did he, Dai, …

DE: Oh no, indeed …

YC: … but he then took all in, and he could sense how people felt at the time. If there was real pressure then he could feel it. And I always like to tell people that, if we had a big problem in the morning and by … at quarter to twelve he would send one of the technicians to, with a piece of paper and pencil, and say 'What would you like for a sandwich?' So that was an indication that you are not going to the Staff House for lunch, you see. And once or twice he came himself with a piece of paper and a pencil, asking the same question. So, I mean, he didn't say 'You're not going for lunch' but it was just so kind and, er, considerate. He knew that we had to eat somehow and he would do things like this. So, er, he would sit in that red arm chair – it was a kind of – the upholstery was in tatters and he just sat there and every morning Joan Hart [Tom's secretary] would come in … and … to get him out to deal with his mail. And he didn't like that a little bit! He would come back within a quarter of an hour. [Laughter] So, that's the sort of person he was.

SL: [08:41]. Yes. Very interesting. Now I wonder whether we could turn a little bit closer to the engineering. Atlas was an asynchronous machine, er, something that perhaps the present-day audience will find hard to understand. So, fault-finding on a

prototype machine that was asynchronous must have presented some problems. The classic expression that signalled doom was the expression 'Lost pre-pulses'. Can you … First of all, Yao, can you talk us through fault-finding on a machine where you didn't know where you were to start, so to speak.

YC: Well, because I didn't have any real experience with a synchronous machine – I was not part of the Mark I or [Ferranti] Mercury, you know – so I started on Atlas without this background. And so I just worked way through with Dai, with Frank [Sumner] and Ianto [Warburton] and Gordon [Haley]. And so in my mind there was only one thing - which was asynchronous. Asynchronous in the sense that we didn't have a clock rate … everything was autonomous. And so you just sort of progressed from one stage to another and, as we called a pre-pulse, we had one pre-pulse and it would go round the machine all the time. Of course the problem came when we lost pre-pulse – you would have heard the term 'We've lost the pre-pulse' and the machine went dead. And then you had to find out exactly where you'd got to - why, you know, a circuit had been broken. And that would take quite a while if you don't really understand the logic … of the system.

SL: Indeed. In trying to get something repetitive that can be observed on an oscilloscope must have been, er, ..

YC: Yes. And then sometime we had to artificially inject something at a certain point and to see whether it would complete the loop …

SL: … how far it would get …

YC: Yes, yes.

SL: Dai, would you have any comments on this, the difficulties of fault-finding?

DE: The problem really which, er, we were really worried about was the fact that when there was a number of items wanting to use the same facility, for example, accessing the store. You would have the processor unit wanting to do something with the store, like get an instruction, you would want the drum which was making autonomous transfers which was coming up fairly rapidly with requests to go to the store, and there was the tape system where we had up to eight tape decks possibly making requests. And so there was the problem that you only had a certain time … you had to decide which device you had to deal with. If it was the drum, there was going to be a crisis time because it was very rapid, the transfer rate. The tape was less rapid and of course the processor could actually be stopped for a time without too much seriousness, so the problem is you had to decide which of these was going to get access to the store.   Now if you imagine them all making a request at the same time then there's no problem because, when you look, everything has settled down and you have this short period when you come to make the decision. But if, for example, during that short period, you were thinking you were going to choose the central processor unit but the drum suddenly came up half way through, then this produced you a partial pulse. And the … we managed to set up an experiment to look at what happened to a flip-flop when you put partial pulses in. And the thing is, it could take a long time to recover. So that all you could actually do was to wait a period of time and in that time what you expected was that most of the time things would be OK but there was always a possibility that it would be, you know, too short a time – you hadn't allowed enough time – and so what we had to ensure was that the mean time between failures

was out in the sort of months region, where other problems arose, and you would be clear, you know; the fact that when this thing was actually wrong would be like just another error. And so, that was really the critical thing as far as I was concerned but, as Yao mentioned, you had to have the ability to inject pre-pulses at a certain time. And of course there was always a comment, that people around to say: 'There must be a bucket full of pre-pulses somewhere, because we've lost so many!' [Laughter]

YC: Of course, the, one of the most important factors is the edge-speed of the pulse because when an Interrupt comes and then you have to make a decision whether that Interrupt should be allowed or not. And it depends on how fast the edge or the speed of the pulse would come along and then … You know, unfortunately in those days sometimes the pulse was really not in good shape. And when you've got a long piece of wire and then the edge-speed can go like this [Yao makes a drooping gesture with his hand] and its very very critical and we had so many problems with that. As I was saying to Dai yesterday, I wish we only had had ten percent of present-day technology – it would have made all the difference.

SL: [14:55] Yes. Now of course you were aware of, um, loading and distance and distributed capacitance and you had carefully-designed cable drivers and things, but even so there must have been occasions when you felt that the pulse shapes were not so good and, at that late design stage, perhaps you couldn't do anything about it, or .. what did you feel?

DE: I never felt that things were on the edge, from that point of view, when the machine was produced. As far as I was concerned, things were as satisfactory as you could make them from an engineering point of view. On the Atlas we dealt with hundred nanosecond pulses, basically. And for the sort of decisions I talked about earlier, there weren't many decision-circuits that you had to provide. And so you could always choose the best possible transistors to make, in that instance, things actually only look for fifty nanoseconds and improve the mean time between failures and so on. So you could actually take a decision to do something special for a limited number of cases. And I think that sort of approach still persists, really.

SL: Right, yes, OK.

YC: There's only … Till my dying day I would still remember the output of one certain long-tailed pair, Dai, because I wanted to find another driver and unfortunately it was way down in the rack, so there was …

SL: Another spare driver …

YC: Yes, a piece of wiring and that edge was really bad. Instead of going down like this [Yao makes a rapid downwards swoop with his hand], it took 25 nanoseconds, you see. And at first I thought it was unique for the Manchester machine. By the look of the London machine [and] on the NIRNS [the Atlas at the National Institute for Research into Nuclear Science (NIRNS) at Chilton, near Harwell] machine they were all the same. I could never understand why that piece of wire should have made all the difference. But thank God, it was not critical, the driver, at that point.

SL: [17:12]. This, er, relates also to faults that might have been hardware or might have been software. There was a grey area often, wasn't there, when the machine packed up. A simple example, I think, was given by Dave Howarth at yesterday's lecture [on 5th December, at the Atlas Symposium] when, during the Sir John Cockcroft demo, the status of the printer changed from 'start' to 'stop' ? Dave suggested, I thought, in the lecture, that was a hardware fault – but it may not have been a fault?

DE: Yes, I must say I didn't myself actually recall it being identified as a teleprinter problem. But, er, he was implying that this was an arbitrary fault which came up on the equipment and from time to time that obviously caused difficulties. Now, I can't really recall engineering faults of that sort of nature which we didn't actually, you know, in the end, be able to do something about.

SL: No…

DE: It could easily have been a fault of the equipment but, I mean, for that sort of equipment you would change it to something else, is the first reaction, to see if it was inherent in the equipment, or some arbitrary thing or whether, you know, it was peculiar to that one.

SL: But it was a standard piece of peripheral equipment, a teleprinter. I just wonder whether those devices, um, if they were not attended, if they were not in use for a certain number of seconds, they automatically switched to a quiescent state that was 'off'.

YC: It's always a very very, er, sensitive issue when you talk to me, talk to my colleagues on the engineering side, or to the software people. And, er, I like to think that the final score was 'fifty-fifty'. [Laughter]. And, I mean, on a number of occasions we had to go through the instructions in the software line by line and sometimes I would say to them: 'That's wrong, because …' One case I remember very clearly was with Robin Kerr on IIA [Intermediate Input A] in the Working Store and they were getting inconsistent answers. And I went through a few instructions and found the compiler did not clear the working area in the Working Store. So, it all depended on what was left in that area when we started executing the program. And Derrick Morris, I had many many talks with him … and you see … But luckily we understood each other, and it was never sort of at loggerheads. We understood each other's approach. And finally, I think, in ninety nine cases out of a hundred we got everything resolved.

SL: There was, I believe, a case with Graham Maclean, Mac, and a program that was running inconsistently. Can you remember … ?

YC: No, not really, I couldn't …

DE: I remember things, not appropriate to Graham Maclean, but to Owen Mills where he was getting one of two results in his program … er … consistently. So he would either get one answer or another answer, right, and this in fact turned out I think to be something in the arithmetic unit …

YC: Yes, yes …

DE: … with the timing that, you know, it had been adjusted too tightly …

YC: Yes …

DE: … for some reason. And that was a problem. And it was producing consistently, in a sense, but you couldn't tell when one result or another result. And Owen Mills was able to demonstrate that. Now equally, there was another problem we had where the machine started behaving badly, no doubt about it, right. And eventually this was traced to the B store. And it was when the machine was actually running, the Atlas was working. And I was called back in because, I mean, being responsible for a lot of the B store, I was called back in to help sort it out.

SL: So this was 1963 by now?

DE: Yes, I should think so, yes. And I was called back in and, er, there was no doubt about it, the thing wasn't behaving to the same standard. And what this was eventually attributed to was a fault which had developed on electrolytic capacitors. But instead of behaving like a capacitor, they suddenly developed a resistive component. And we had to change - I mean, it took us, you know, a day or so to discover that, you know, it was all the twenty four amplifiers that we had to deal with, and we had to change all these electrolytic capacitors on the [printed-circuit] boards. So we lost about three days of Atlas time overall, by the time we located it and by the time it was then all corrected, before it could be used again. But, I mean, that time we not only worked all day, we worked all night as well to solve the problem. But there was no doubt that was an engineering problem, and it was a faulty component.

YC: And the other case was Jodrell Bank, with John Davies. And then we started to use his program as a test program, you see. It came out with consistently inconsistent results and it was once again a problem with the accumulator, with the floating-point arithmetic unit. And, but you see, of course some faults were due to just intermittent contact problem.

SL: Yes, electro-mechanical contact problems ..

YC: We had seen the printed-circuit boards …

SL: … yes …

YC: … and all the joints were done by hand with soldering iron. I mean, it's just incredible when, when, you look back.

SL: Yes, five thousand printed-circuit boards, all done by hand …

YC: … all the joints. I mean, how many soldered joints are there on each board …

DE: And, I mean, the problem was, it wasn't just simple wiring. I mean, in a sense there was, the short wiring was all a simple connection but over a certain distance it had to be twisted-pair wiring and, of course, on longer distances it had to be coaxial cable. So, you know, the process in itself of just making all those connections was quite tedious. You had

to be incredibly accurate. And it involved doing that job of connecting and then checking that it had been done correctly, and so on.

SL: Yes.

DE: I do remember one fault on Atlas where the machine had been working for some time and where, for example, er, when we found the fault that was occurring, it suddenly developed. The pin on the back of the socket had never been soldered. And we couldn't even find a wire that might have been twisted around it, you know. There was not a connection to it at all.

SL: So, there never had been ….

DE: How on earth the machine had ever operated we never really explained. But we had to, then, re-solder a wire and connect it to the point it should have been to, and then it was OK. So, you know, a most peculiar error which we never explained.

SL: Right. It only made itself apparent after the machine had been working for some time …

DE: … for some time.

YC: And the sort of conditions – the temperature, humidity and everything … I mean, we had so many boards which we took out and sent to the factory and they came back with "NFF" – "No fault found"…

SL: Right…

YC: … What do you do? When the next fault came along, they would put the same board in and it worked. So, you know, it was very difficult … No wonder a certain high-level senior manager in ICL said: "Atlas is un-maintainable".

SL: Right …

YC: I wouldn't like to name that person, and I was furious when, at an internal meeting ….

SL: … yes, yes …

DE: It's quite interesting because one of the people who worked on the London Atlas said, you know, initially when the London Atlas was in, it was frantically interesting and they were very busy but then by the time it had been in about three years he was bored stiff because there was only really one sort of fault a day, and there was nothing really for him to do …

SL: A maintenance engineer?

DE: Yes.

SL: [26:25]. Yes…. Yes.   Let us turn to wider issues now and things that were going on in the rest of the computer world. Within the company, within Ferranti, there was the Orion project and we might mention that in a little while. But of course IBM was the, became the main challenge and STRETCH was, er, came into being. Did you have any, any contact whatsoever with STRETCH developments with, er, their circuitry, their problems? Was there any influence at all between the IBM group and the Ferranti Atlas Manchester group?

YC: I was totally ignorant in those days, I mean… Dai has got a lot to say …

DE: I went on a three-week visit to the States in 1958, right. I visited IBM at Poughkeepsie. What amazed me was that there were groups working there who didn't … I went to see various different groups … and they didn't know what the other group was doing …

SL: … right, within IBM …

DE: … within IBM. That was the first thing I learned. The other thing is that I talked to them about, for example, the fast carry circuit and what we were trying to achieve. Er, we'd only done, you know, limited experiments at that stage but it was looking hopeful. But they were quite unbelieving, for example, that this was a possibility. And, as for STRETCH, I mean, when they made publications, those eventually came through and we were able to see them. But there was no real contact with any IBM engineers, you know, or things like that. In the early days, people like Samuel [Arthur Samuel (1901 – 1990) of IBM], who had been involved with the 701, used to come regularly to Manchester. And he was very interested when we had the floating … we showed him the floating-point Meg [prototype for the Ferranti Mercury computer] working, which was the predecessor to Atlas, you know. About three months after he came looking at the Meg, they [IBM] stopped the production of the 701 with which he'd been critically involved, and they came out with the 704 which had floating point [hardware]. Now I'm not saying that we had any influence, but he was impressed with the floating point. And I'm sure they'd been working on it but they did come out relatively quickly with a change in machine even. Now, there was never anything the other way round, really. We didn't have regular visits to IBM or see things that were done. I'd made an arbitrary visit - er, I came back – I think I came back with a few, for example, three, a piece of ferrite with three holes in that they were trying to make a storage system with. So, they did give us a few things like that, you know, but there was never any real contact at circuit level, or detailed level, of any logic in any way. We just knew what their aims were, and we knew what our aims were.

SL: Right.

YC: I can remember one thing Dai pointed out, when he came back, about STRETCH: 'Never get too far ahead of yourself in the pipeline'….

SL & DE: Yes …

YC: … and we really remembered this all the time, every time we wanted to do something … we said 'No, no, no, that's too far'. You know, a jump goes the other way, you've got such a lot to unravel.

9

SL: And of course, in the event that was the weak point of STRETCH.

DE & YC: Yes, yes.

YC: … and that was the message that Dai gave us. [Laughter]

SL: [30:24]. Now, on the device side – transistors – in the early days – I suppose we're talking about 1958 – the emphasis was to try to make a high-speed parallel adder. Now, was it the case that there were just no suitable transistors manufactured in the UK? Your supplies were coming from the States: how did you obtain these things, which were quite expensive? Tell us a little bit about device …

DE: Yes. The surface barrier [transistor, the SB240] had really been announced by Philco, who were the manufacturers. Now they'd not only announced the transistor but they announced they were making a machine [ie a computer, the TRANSAC]. Philco were making a machine with it. And one of the interesting things to us at the time was the fact that these were called *direct coupled* circuits. So you could actually connect the collector of one circuit with a straight wire to the base of the next. And the transistor was so accurately defined in terms of voltages that this actually worked. But the cost of the transistor, due to its manufacturing process, was very expensive. The early …

SL: Can you give us some idea, compared with the price of a weekly salary or something …?

DE: Well, I think in those days the first few transistors we got were about £50 each.

SL: In today's money that would be £500, would it?

DE: Certainly, I would think.

YC: A lot more. I was getting £50 … I was getting £20 a week, so you can imagine …

SL: So, two weeks' salary, at least …

YC: … as a graduate, £20 a week …

DE: Yes. Now, so there was no doubt that we couldn't think of, you know, a large number of transistors costing that sort of money. There was no way we could make a machine … Philco, you know, since they were the manufacturers of the things [transistors], they could get round that problem. But there was no way we could think about it but we did want to think about using them perhaps for special purposes, like for making this, er, *looking time* [decision circuit] to decide which device we were taking, and so on …

SL: … decision time …

DE: Yes, we could think about it in that. And of course then came up the idea of 'could we think of it just as a switch itself', so that, when it was a switch, it stayed as a switch and it would transmit a signal straight through it …

SL: Like a relay, a mechanical relay?

DE: Yes. So we could think of it like that. But we went to Mullard [a UK manufacturer of transistors] – I remember paying a visit to Mullard down south at the time, to discuss with them what the possibilities were for transistors. And we did mention the comment I showed on the slide [at the Atlas Symposium on 5th December, the previous day]: 'First into the Future'. And I think we did mention that to Mullard, and they said to us: "We are very happy to be *third* into the future, because it is really too expensive to be anywhere near, like, first". And they said to us: "We are making some fast transistors but you need to maintain voltage across them". And these were called *graded base* transistors, which had a sort of built-in field within the transistor to accelerate the carriers across. But they said, in order to do that, you need three or four volts across it at all times.

SL: Yes, you mustn't let it get into saturation …

DE: So, switch it on, switch it off. Don't ever operate it lower [collector-base] voltage that that, because the grading won't work in that connection and it would stay a long time in the saturated state. And you'll be a long time switching it off. And that was true. But, of course, the cost of that [Mullard transistor] was, I think, well under a pound.

SL: Per transistor?

DE: Per transistor. So, you know, it was a feasible device. And therefore we had to think in terms of a logic circuit for general use which was in that category.

SL: [34:31]. Yes. What about a similar discussion of core stores, ferrite cores? You presumably looked at available British manufacturers, and American manufacturers?

DE: Well, we actually did some work ourselves in that area, in conjunction with Plessey [a UK company], who supplied us with sample layouts of cores, as we required. And at that time, 1958, was the time that Mercury was being delivered and that [ie Mercury's core store] had a ten microsecond cycle time. Our initial thoughts were to look at partial flux switching, which was where you didn't switch the core from one [stable] state to another but like, you say, could we switch it to 20% or 30% and achieve a speed improvement by that factor? But we were worried about the stability of that sort of temporary position – part way in the flux switching. And what we found, we seemed to find, that in order to get stability you needed two cores per bit….

SL: Right …

DE: … in the system. So we had at least discovered that. But we were also, er … In the core manufacturing, they [Plessey] were also thinking of moving to smaller cores. So, whereas the Mercury cores were 80 thou diameter, the next core under construction was 50 thou. And I think Plessey were telling us that, you know, they were thinking in terms of coincident-current storage with a 50 thou core and that might approach the two microseconds [cycle time] that we were really thinking we might like to achieve on, you know, for the main RAM of Atlas.

SL: At this time it sounds, from what you were saying, this is 1958, as if the university group, the academics, were acting as the R & D department of Plessey. Was this in any

way collaborative, or were you pressing ahead further than the Plessey engineers themselves?

DE: Yes, I mean, at that stage, because we actually, I think there were papers published in 1959 in the IEE [Institution of Electrical Engineers] concerned with this. And I think we were actually - we weren't actually thinking of the smaller cores, we were just saying to Plessey 'are you going to make some smaller cores?' But there was a problem with getting wires through smaller holes, and all this sort of argument going on. But the partial flux switching was really being done at the university, as an independent action. Independent of Plessey, although we were talking to Plessey about it, talking to the people we knew there, and they were supplying us, I think, almost free of charge with, you know, small amounts of cores: sort of one small plane, and things like that. So I'm sure that, at the time, I had about half a dozen different versions of things that we were looking at.

SL: [37:26]. And at this time, Yao, you were mainly on arithmetic …

YC: … on the B and C, as we called it …

SL: Yes …

YC: … doing B logic and also the central control. But I spent many happy hours with Mike Lanigan working on the B store and working on the [main] store. And Mike was a very very dear friend – and unfortunately he's no longer with us.

SL: Indeed. And also round this time Frank Sumner, who was by training a Chemist, and I guess came into the group as a programmer … Er, it seems to me he got more and more involved in the logical design. Talk to us a little bit about Frank.

YC: Well, Frank is still a very very good and close friend and, poor Frank, he was smoked out by Gordon Haley and myself! We spent many happy hours in Dick Grimsdale's office, Dai, if you remember, and the two of us smoking pipes and Frank was sitting in the middle … And, yes, Frank made a tremendous contribution and … Because we tended to get very blinkered, if you like, and he would point out 'No, you can't do this'. And once again Dai's advice came in: 'Don't go too far ahead of yourself'.

SL: Yes.

YC: And, you see, so, yes, and Frank made an … Of course, he did the *Drum learning program* and also the test programs. He wrote a lot of test programs for the whole machine, you see …

SL: So, he was seeing a slightly wider picture …

YC: Yes, yes, definitely. And I think he got a very, as good a concept of the machine as Tom and Dai…

SL: Yes, yes …

YC: … so as I was saying to Dai: in my mind there were three people: Tom, Dai and Frank, as the three, you know, leaders.

SL: [39:36]. Yes. Now could I switch again to the broader picture for a while and talk about the University of Cambridge. They were given, as we know, some central Atlas hardware and more or less told to get on with the design of a simpler system. I was intrigued to hear that they worked on Tunnel Diode caches for the machine. Was there any interest in, or work on, Tunnel Diode stores in the University of Manchester? What did you think when you heard that this was going on?

DE: Oh, I mean, we were well aware of tunnel diodes and I think Mike Lanigan actually wrote a paper on the possibility of a tunnel diode store. And we certainly mentioned it. But they were again a limited supply available and limited reliability. So … Plus the fact, really, that we were convinced that the *One level Storage* system was the right answer. And in a sense we were against the use of, sort of, small buffers for operand and instruction access, you know. Whereas, I think, you know, because of the supply situation and things, Cambridge were looking at this alternative, which was a bit cheaper – perhaps not quite as cheap as they made it out to be. But it was appealing to Ferranti because, even when we were making Atlas, there was a group at Ferranti, I particularly mention Gordon Scarrott, for example, who kept saying, you know, 'Do we really want to go ahead with this *One level store*? Wouldn't it be better to use a buffer technique?' And so on. So I think there was a group at Ferranti who were, if you like, thinking the same way as Cambridge.

SL: … and were a bit sceptical about …

DE: Right. Yes.

YC: Of course, Gordon [Scarrott] was also a Cambridge man himself. [Laughter].

SL: [41:49]. Well, and there must have been a certain amount of rivalry, if that's the right word, between the Orion group and the Atlas group. Two large machine projects more or less running concurrently in the same, er, in the same company. First of all Yao, and then Dai. Yao: you were by that time working for Ferranti. What …In terms of engineering facilities and management structure, was your life an easy one or did you find some clashes?

YC: Oh, yes, definitely. I mean, it was very clear there were two camps at West Gorton. You've got Orion and Atlas, and, er, … But, you know, when we were all at the university, the difference was not so apparent. But when I took the prototype of Atlas …

SL: … the Pilot version?

YC: … yes, the Pilot version … to West Gorton while the computer room [at the university] was being prepared for the final production [system] and I had to re-commission the thing [the Pilot version, at West Gorton] and I would work till quite late. And after, I would go in in the morning at nine o'clock and it was not working any more, you see. And it took some time to find out what was wrong and in the end I found that some boards, you know, had been changed. Because there were certain boards compatible to both machines on the peripheral side …

SL: Yes.

YC: And I could never understand it. I said, 'it was never like this'. And I just put in new boards and, er … Until five years ago, five years ago, when I saw Virgilio Pasquale [ex-Ferranti West Gorton] and he was responsible for the Orion machine in the lab at the time. And he confessed to me: it was *him* who took boards out of my machine … So, you could see the kind of thing … But, on the whole, I thing we basically had respect for each other. I mean, the Orion was a good machine; the only problem was the hardware was not up to it because they tried to incorporate so many facilities by means of hardware and the hardware was not up to it. The logic was not up to it.

SL: [44:30]. Do you want to make any comment, Dai?

DE: Yes, I'd like to make some comments because, really, it all stems from the first [joint University/Ferranti] Progress Meeting, my contact with Orion, so to speak. Because at that meeting Tom had to say, er, what MUSE was doing. This was 9th November 1959, the first Progress Meeting. So Tom was reporting what the state of MUSE was at that stage, because MUSE was changing into Atlas. And so, that's the Progress Meeting at which we talked about the Pilot model of the Atlas, rather than MUSE. And so Tom was reporting where we'd got to, so to speak, on MUSE, right? So he was reporting on the Fixed Store particularly, on the parallel adder particularly, OK, and Keith Lonsdale [Ferranti West Gorton] and Gordon Scarrott were at this meeting and Keith Lonsdale was saying: 'What's the state, what should the university expect in terms of peripheral equipment?' because building the central processor was our main action. We weren't thinking in terms of peripheral equipment, but obviously from the atomic energy people's point of view and NRDC's point of view, they wanted a much more comprehensive approach including many more, and many improved, peripheral equipments. And, in a sense, er, Keith Lonsdale was conveying to us the fact that he wasn't expecting anything from Atlas on peripheral equipment, right, that that would all be done by the Orion people, certainly at … because this was the initial sort of indication of how things were going to go and that Atlas could expect peripheral equipments to be provided similar to Orion. Er, in this … At that meeting it was decided that the Pilot model would have drum store and magnetic tape store attached to it, right. Now in fact that never happened, because of problems that arose, in a sense. But at that first meeting, that was the original intention, OK? Now …

SL: So the problems were?

DE: The statement was, for example, Orion was using a drum called the MD4. I think either at that meeting we said we would be liking a drum with faster access and I think people like Gordon Scarrott said: 'Well, that's no problem. We can supply, we can supply a drum with faster access – it might be called the MU5 – er, the MD5'. OK. And that was just accepted. But at that stage there was no indication precisely of what the electronics would be for the drum. Similarly, later on, it was announced that Ferranti Edinburgh were going to be brought in, to take decisions on the magnetic tape. And, er, because Ampex were a big manufacturer …

SL: … American manufacturer …

DE: … and of course people at AERE were looking for advanced performance of the magnetic tape and there was a sort of move from half-inch tape to one-inch tape. And some of the tapes being used were metallic rather than plastic, for example. And so it was

a state of flux on the magnetic tape. So Edinburgh were given the job of choosing a tape deck and doing this. And later on, like, you know, months later, like August, they said they'd received a tape deck from Ampex, they'd investigated it, but it was unsatisfactory in its acceleration rate, I think in its final speed, and they didn't like the read- and write-electronics that were supplied with it. And so they talked to Ampex and the machine, the, Ampex had said, well, our TM2, which was in fact the deck that was later used on Orion and us, was going to be a year, was going to take a year to actually get the first model back to us. And so, you can see now why it wasn't connected to the proto, to the Atlas Pilot model. So, there was that situation. And, for example, er, but Ferranti Edinburgh were to do all of it. Now, er, and it was going to be common to Orion and Atlas. Now, what we found was, later on, Edinburgh came back and said: 'Look, we have the production facilities and the test facilities. What we're actually short of is design capability to finish off the design of the electronics which we're now doing ourselves'. And in fact Dave Aspinall [from Manchester University] and I went to Edinburgh for a week. They asked if Atlas could help, right, and we said we would go for some time to sort it out. And we went. I think we turned up at eight o'clock every morning for a solid week, right, and we worked till late at night, and we managed to get basically it done. I think there was a fellow called Sam Alexander who we were talking to at Edinburgh, and he approved it for both Atlas and Orion, what we'd done. And I think that was later approved by Gordon Scarrott as well, and so, on that side, it took a relatively short time to sort out the little problem and the shortage of effort that had occurred. And after that, it went reasonably straightforward.

SL: The impression I get is that, although the nominal responsibility for both drums and tapes was the Orion teams', that when it actually came to, with the tape decks, getting the right performance, it was the Atlas people who made the difference and actually produced the circuitry that was up to scratch.

DE: Yes. And I would say that, for the tape, it was not too much of a problem, right – it was only a fortnight of two peoples' effort. But with the drum, where I worked with Eric Dunstan [from Manchester University], for example, from the early days we had problems. These were problems of, really, to do with the construction, the mechanical construction, of the drum. And the fact that you were using block heads. Now there were eighteen heads in a block and what this meant is, when you made the heads, you had to get most of the eighteen working. So it wasn't like you could be assembling heads one at a time, and throwing the odd one away if it didn't work. You had to get the whole block working. And the fact there was a whole block having to be put very close to the thing [to the drum surface] meant that you, you had to align it mechanically and physically in that position. There was no adjustment. And so if, over that sort of six-inch area of the block head, there was some change due to either temperature, right, or centrifugal force, then this was a worry. And in fact what we observed when – we decided to run 25 tracks in parallel, right, and you couldn't run these heads that way, because there was too much cross-talk. So that was a problem. The fact is that, over that distance, the read signal was different because of the different head separation due to temperature. And so we did work very hard, and Gordon Scarrott kept assuring us that these problems would be resolved. Extra holes were put in the outer casing to get proper ventilation. The ventilation for the drum was kept on all the time, so that the inside and the outside would cool at the same rate, hopefully, and all this sort of thing. But, looking back at it now with hindsight, I wonder that we were patient for so long. Because I think it really took until May 1963, when the AERE people complained about the reliability of the drums when they were using our Atlas …

SL: … the MD5 drums?

DE: … yes, that a decision to change the drum was made.

SL: To Bryant?

DE: Yes, and the, er, London Atlas was delivered with the old drums but then updated as the improved drum came into being. I think there was also a problem we had, er…. Initially there was a drum team on Orion. But I think the people disappeared, and we had quite a lot of discussions about the precise selection and arrangement techniques for doing this and I think Eric and I proposed at one time that we change from just a single coil head to a centre-tapped head. And this eventually got approved of. But in the end I had to say to Gordon Scarrott: 'The system we're going to build is this. We've had three or four gos, at three month intervals, to agree with your system, you know, this is it. Now we can't go on any longer. We have to build it'. And that's in fact what happened. So … And he had to admit at that stage there was really no experienced [Ferranti] engineer left on the drum, right ..

SL: … on Orion at West Gorton …

DE: Yes. But the Orion drum was, as far as I know, was always the MD4 but we had the extra, you know, half as much speed again, to get the improved access time, and of course heads in parallel. So we were actually, we were actually getting faster rate of digits off the drum. Er, because, in the end, what they were proposing, because of this variation between heads, they were proposing that we had automatic gain control. Well, in our case we'd have had to have 25 automatic gain controls, which we didn't want to put in.

SL: [54:21]. Can I, er … This is most interesting but I realise we've gone on quite long. Can I, as it were, wind up these most interesting conversations with a little bit of what you might call 'family and social' matters? Both of you, and all members of the team, were obliged to work long hours and, as far as I can detect, were more than willing and happy to work these long hours. What about, what about friends and family? How did, how did others feel about this? Yao, what would you say?

YC: Well, I think my family was very understanding and I never had any real problem. Certainly, only thing I would make sure was that I would take the girls to school every morning before I came to work. And I would go home most days for dinner in the evening, give the girls a bath, put them to bed, and I would go back to work. So, they were happy with this arrangement. But at the weekends, no problem. So, on the whole, I was very lucky.

SL: And what about you, Dai?

DE: Yes, well, I mean I had two children - one in 1956 and one in 1958 again. So, you know, my wife fortunately was, er, you know, being at home at that stage. She was not actually working in those early days, so I was very fortunate. And I did the same as Yao. I mean, you know, the children went to school on a morning, er, my elder child had to be taken to school and things, but we were very lucky in having a grandma as well who could actually do some of that task. So I was very fortunate to be relieved of quite a lot of that

activity and, er…. My next-door neighbour, for example, didn't believe that I worked at the university because I sometimes worked nights, which he couldn't understand at all because he worked as a bank manager. And I didn't seem to have all of August off, either, you know. I just took a fortnight's holiday, and that was it. So he found that very surprising. But I think from a general point of view I was very fortunate because in 1959 Tom and I were in, er, went to the University of Pisa as a result of an invitation from them and we spent a week at Pisa University discussing Atlas. And in 1961 we were both invited by the Academy of Sciences in Moscow, where we went for a week. And, in fact, after the week they invited us to go to the Black Sea Resort, but, so, it was very nice, but our Visa was only for that limited time and we just said 'I'm sorry but, you know, we can't go because of the Visa', that was our initial reaction. And then they said, 'Oh, don't worry about the Visa, we'll take care of it'. But we were a bit worried about that and we said then 'Well, we really have to get back to the university [of Manchester] because of our teaching commitments'. So they had to leave us go on that basis. But I must say we were a bit worried, but it was a generous gesture, so to speak. So Atlas, you know, apart from us having to build it, it was also, you know, it also gave us a social contact with other people. And being invited to give lectures and so on, externally, which we did.

SL: Right. Well, I can hear singing noises off. For the record, we're making this recording by kind permission of the Students' Union, who have loaned us their most up-to-date digital studio. So, Yao and Dai, let's end this and may I say thank you very much for a most interesting conversation.

DE: Thank you very much, Simon.

YC: Yes, Simon, thank you very much indeed.

Elapsed time at end of recording: 58:24.

MP3 file size = 137.2 Mbytes. Recorded on 6th December 2012 at the Fuse FM Radio studio in the Students' Union, University of Manchester, M13 9PL.

Biographical notes.

D B G Edwards.

Born	Tonteg, Nr. Pontypridd, SouthWales, on 14th March 1928.
Education	B.Sc. Physics, University of Manchester, 1948. M.Sc., Electrical Engineering, University of Manchester, 1949. Ph.D., Computer Engineering, University of Manchester, 1954.
Employment	University of Manchester: Assistant Lecturer, 1949 – 1953;

Lecturer, 1954 – 1959; Senior Lecturer, 1959 – 1964; Reader, 1964 – 1968; Professor of Computer Engineering (later changed to ICL Professor of Computer Engineering), 1968 – 1988; Head of the Department of Computer Science 1980 – 1987. Retired in 1988; Professor Emeritus, 1988 to present.

Brief summary of involvement with computers, especially with Atlas.

From 1948 Dai Edwards worked with F C Williams and Tom Kilburn and others on the Manchester Mark I computer, and then on the transfer of the design to Ferranti Ltd. The production version, the Ferranti Mark I, was first delivered in 1951. Dai then worked successively on four more University of Manchester computers, namely the Meg, MUSE (which became the Ferranti Atlas), MU5 and MU6. Three of these machines had their industrial derivatives – the Meg becoming the Ferranti Mercury and MU5 contributing to the design specification of ICL's 2900 range. On the MUSE/Atlas project and the MU5 project, Dai led the hardware side of the teams.

Dai has been the inventor or co-inventor of more than a dozen computer patents and has published 45 significant papers on computer-related hardware. He has acted as a consultant to a number of companies, including Ferranti Ltd. (from 1950 to 1963) and ICT/ICL (from 1963 to 1972).

E C Y Chen.

Born	Shanghai, China on 16 August 1929.
Education	B.Sc. Electrical Engineering, University of Manchester, 1957. Post-graduate research in Switching Characteristics of High-Speed Semi-conductor Devices, Department of Electrical Engineering, University of Manchester, 1957-1959 under :Dr. D.B.G. Edwards.
Employment	Ferranti Ltd./International Computers and Tabulators Ltd. (ICT), Computer Department, Development Engineer, 1959 – 1964; University of Manchester, lecturer in Computer Science, 1964 – 1967; ICT/International Computers Limited (ICL), Systems Manager/Project Manager, 1968 – 1982. World Systems Limited, Systems Consultant assigned to the Statistical Office of the European Commission (Eurostat) and ICL (Germany), 1982-1989; World Systems (Europe) Limited, Managing Director, 1989-2006.

Brief summary of involvement with computers, especially with Atlas.
Yao Chen worked within Tom Kilburn's group on the MUSE/Atlas project as a hardware designer from September 1957 until 1963, when he progressed to the role of principal support engineer for the three Atlas installations (Manchester, London and Chilton). In October 1964, he joined the Department of Computer Science at University of Manchester as lecturer, responsible for teaching undergraduates on computer system design and supervising electronic laboratory work.

In 1968, he took up a post of Regional Systems Executive in the Universities and National Research Region of ICT (later ICL) and progressed to management appointments as Systems/Senior Technical Manager in the Government and Public Sector Sales Organisation of ICL, managing technical support and later large-scale projects for ICL in Germany and Luxembourg until 1982. He worked on assignment from World Systems Limited as a Systems Consultant to the Statistical Office of the European Commission (Eurostat) in Luxembourg, advising on Informatics Management between 1982-1989 overlapping with project management tasks for ICL Germany between 1985 and 1989. He founded World Systems (Europe) Limited in 1989 and managed a large number of projects for statistical data processing for Eurostat, until his retirement in 2006.

Yao is a co-inventor on one of the Atlas patents. He co-authored a paper on the Central Control Unit of the Atlas computer with G. Haley and F.H. Summer and a paper on Test Programmes for the Atlas computer with A.B. Cooper, presented at the IFIPS Conference 1962 in Munich and IFIPS Conference 1965 in New York respectively.

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 that 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   *23rd May 2013.*
lavis@essex.ac.uk