AICSCSHistory_Interview008_GeorgeRoss

Vassilis Galanos 00:02
Excellent. So today is the fourth of July 2024. My name is Dr. Vassilis Galanos, and I'm here with Dr. George Ross. This is part of the Edinburgh AI and CS History Project. And so before we start delving into the specifics of this interview, I'd like to ask you, George to briefly describe your background, your contributions to your field in computer science, and perhaps what brought you to Edinburgh University in the first place. When did you arrive here?

George Ross 00:42
Okay, I actually came into computer science slightly by accident, because I'd originally come here [to Edinburgh University] with the view to doing mathematical physics. I'd done a bit of programming in sixth year studies at school. So the sort of thing, you write your Algol 60 program out on a coding sheet, and the teacher posts it off to Dundee, and then they do something and run it and about a week later a pile of paper arrives on your desk, and program has either worked or it hasn't. And then you send it around again, and a few iterations later, you've done your program, which is all well and good. But I remember wanting actually to calculate pi, I think it was [rather than just some textbook exercise, that is]. They must have wondered what was happening because they typed the program in and ran it. And then it must be a minute later on somebody obviously pressed the abort button, because nothing was coming out in the printer. So they sent it back and said it didn't work. And somebody had obviously looked at it and thought oh, it actually should have worked. So they stuck a print into the into the loop and ran it again. And they must have pressed abort pretty quickly because they only got about two inches of printed paper as the thing gradually converged on pi. So anyway, we did that kind of thing. But most of the numerical stuff was on the hand cranked calculator things. So anyway, I came to do mathematical physics. And this is 1974, I went along to my Director of Studies. They said, okay, you're doing mathematical physics, you want to do Maths 1A, Physics 1A, and Applied Maths 1, nice and straightforward. And then he must have looked at what I'd done in the past, well, this is sort of paraphrasing slightly. He said, you're going to be bored, you've seen all this applied math stuff before. Why don't I sign you up for Computer Science 1 as well? So he signed me up for Computer Science 1, just as a kind of extra course. And I did my first year in Physics and computer science along the way. And physics wasn't actually that that interesting, I would say. Computer Science had better toys, so it was more interesting. I kind of thought after having done first year, I really wanted to be a statistician. So I did my two years of Maths, Applied Maths or Math Statistics and Computer Science, and then did Maths and Stats as my first degree. And at the end of that I was looking for jobs, as you do. And I can't remember where it was I applied for a job as a statistician at one of the research institutes. And they offered me a job as their system manager for their computing system. So I thought that was a bit strange. Anyway, I was wandering around JCMB again. Oh this must have been after my MSc, because I did a BSc in Maths and Stats and then a MSc in Statistics. And I must have been wandering around JCMB and happened to see this notice on the end of one of the corridors, saying that Malcolm Atkinson had a studentship, anybody interested? Well, I went along and knocked on his door. And because I'd done the two years of CS, they knew me. So Malcolm said yeah, come on, you can have the studentship and do the PhD. So I did the PhD in database systems really, looking at what are called differential files. It was a way of storing various different versions of a database in a file, trying to work out whether it could be done efficiently. So that's pretty much how I ended up in computer science. And

then at the end of the three and a half years Malcolm was looking for somebody to do work on one of his projects. So in 1982 I got employed by the University, and in various forms after that I was employed by the University for the next 40 years. First of all for Malcolm's project for a wee while, and then I'll come on to this in a bit, writing a filestore system, revised version of it, managing Sun Microsystems things, managing the network, designing the network, managing Linux, more network stuff, all that kind of thing as it went on, until 40 years later, as I say, I retired. So yeah, the CS was, I think it was more interestingly done. First year Peter Schofield, Frank Stacey and Rosemary Candlin were the teaching staff. And they were actually all really good lecturers. They made the stuff quite clear and sensible and comprehensible. Probably partly why I was kind of hooked in was that the lecturers were good. I suppose the physics ones were probably good as well, but I like the CS ones better. And second year was quite an interesting year because they had quite an interesting bunch of lecturers. Hamish Dewar was the, I think the course coordinator, but he was teaching on it. John Gray was doing some hardware stuff. David Rees was doing compilers, I think. And I'm pretty sure Robin and Rod were doing theoretical things on the course, Milner and Burstall, yes. So it was actually a quite a high-powered team that they put on to second year's computer science. And it was an interesting course. Stats was interesting too, they had an interesting bunch of lecturers from stats. Maths lectures were a little bit more variable, but the stats lecturers were generally pretty good, which is why I kind of carried on doing statistics, and then the MSc in statistics and then kind of merged back into CS. I suppose my stats MSc project might be kind of machine learning in a way, “discrimination methods in the diagnosis of breast cancer”. They gave me 1174 sample records and said can you do as well as the clinician? So I got about a dozen different variables from the records, and then a diagnosis at the end. And so I was running standard statistical regressions and things against it. And then I came across what amounted to a machine learning one, it was Jackknife likelihood or something like that, and ran that. And that was actually best of the lot. And it came up just about as well as the clinicians were doing at that point. So it was early machine learning, but it never went anywhere, because I went off into databases instead. So that was how I ended up in CS in the first place.

Vassilis Galanos 06:38
Interesting. I'm curious to find out more about the tests you might have done in terms of diagnosing through that early form of machine learning. Today, people are very much skeptical about machine learning in medical diagnosis, because of the false positives or false negatives. What was the perception then?

George Ross 09:14
Yeah, well, I've scanned my MSc thesis, so I can bundle it up and forward it to you if you like. If it makes interesting reading or not, I don't know. I was running it on an EMAS system which was a little bit heavily loaded. And so I was writing things in assembly language, machine code, to try and optimize the speed of these things. I think it was probably something that somebody who had done CS would be able to do, but probably not anybody who had just come in from the statistics side, or the maths side. They wouldn't have necessarily been able to optimize it to make it work on the machines we had at the time. Because this was on EMAS on the ICL boxes, so two EMAS machines. I think one had a megabyte of store, some discs and drums. And the one that students were landed on had three-quarters of a megabyte of store, discs and probably drums. And it ran quite nicely up to about 32 users, and then it started to get a bit slower.

Vassilis Galanos 10:33
So what year did you start your PhD?

George Ross 10:35
PhD started in 1979.

Vassilis Galanos 10:40
So you lived through a rather tumultuous period in Edinburgh in terms of computer science and AI and the formation of different departments, I guess. Have you got any memories?

George Ross 10:54
I don't have any particularly strong memories of AI to start with, because I actually didn't have much dealing with the AI. I suspect if I'd come and done the CS course, as a straight CS course, I'd probably have done AI is an outside subject. But because I'd come in on the maths side, statistics side, I never actually did any AI as such. It was all on the straight CS side. Yes, it [CS] was basically I suppose quite self contained right at the beginning there, up until about, when did the Division of Informatics form, that was Mike Fourman's stuff, wasn't it? I can't remember the date of that. I do remember being at a meeting up in the Jim Howe room, it must have been, in Forrest Hill, where the formation of, it might have been the Division of informatics it was called at that point, was being mooted, and there was a presentation by various people to say, this is what it's going to look like. And I think the meeting more or less agreed, but I wasn't really part of the prior discussions of it. It was more of a presentation saying this is what it's going to be, we can't see any viable alternative, I think is what it amounted to. So I suspect there'd been quite a lot of discussions going on beforehand, to make it work. But CS itself before that, everything in the 70s was quite small departments. So CS must have been 10 or a dozen lecturers, something like that. One professor, a couple of senior lecturers or readers and a few lecturers. And not that many PhD students. I think there must have been about a dozen of us give or take. Not the kind of factory we've got now, where they're churned out by the 100s. But there must have been about a dozen of us, I think, that kind of number.

Vassilis Galanos 13:21
Do you think Edinburgh's computer science was in some way unique in contrast to other departments around the UK or globally?

George Ross 13:33
I suppose the big thing that Edinburgh had going was they were writing EMAS, because they had been part of the Regional Computing Centre. And they had an ICL machine to run that on. And I think they just didn't like the system it came with and decided to write EMAS on it. And that was, for its time, quite an advanced system. For example, I don't suppose universities around the place at the time, many of them wouldn't have had a terminal room for first year students, for example. So there was a room in the corner of Appleton tower which had 15, maybe, teletypes in it, maybe 20. And the students, first and second year, mostly first-year would go along there and do their programming there. So I don't know how unique it would be in Britain. Various other places, Kent started running EMAS a bit later. But I suspect, other than EMAS, I'm not sure how unique it would have been from the computer science point of view. Later on they went quite heavily into VLSI design, and that was probably a bit more unusual.

Vassilis Galanos 15:05
very large system integration.

George Ross 15:07
Yes. And there were several people were quite heavily involved in that. And, in fact, it spun out a company called Lattice Logic which, I don't know whether they got bought over, it turned into something else. But there was a spin out anyway, John Gray and Irene Buchanan and Peter Robertson.

Vassilis Galanos 15:30
What period was that?

George Ross 15:32
That was after I finished my PhD, so that must have been the 80s.

Vassilis Galanos 15:40
Was that the period of the so-called Silicon Glen?

George Ross 15:43
Yeah, yeah.

Vassilis Galanos 15:44
Do you have memories of that? Like, the feeling of that? I've only read a little bit.

George Ross 15:49
Yeah, I remember people were kind of excited about being able to design their own chips and send them away, and actually get them back as fabricated devices. That seemed to work quite well. And Gordon Brebner spun out, he did it with the programmable logic array type stuff. Yeah, I can't remember what it [his spinout] was called. He went off to America to run that. I was kind of involved in keeping the keeping the ship running, if you like. Not so much in the actual research side of it but making sure that the infrastructure worked. People were buying lots of things like Sun Microsystems machines, and not thinking about how they'd connect them up, and not realizing how much of a network load they were going to be. So I was juggling to make sure that the infrastructure actually all worked. And probably the easiest way to find out about that is I did write a document about networking in computer science and AI. And that's on the history site. [https://history.dcs.ed.ac.uk/archive/docs/dcs-inf-network-history.pdf] That kind of gives a reasonably good description of the state of the network for CS and then for Informatics, from ‘74 on to when I retired, about a year and a half ago. It's probably easier for people to read the document with the diagrams than it is to try and describe in detail without diagrams how it all fitted together, and some of the problems that we came across and tried to solve.

Vassilis Galanos 16:46
So technical problems?

George Ross 18:09
Technical problems, like lack of bandwidth, basically. And if you string a bit of coax cable through a lot of offices, then people are going to catch it with their feet and break it. And this was when you were still allowed to bring your bicycle into the building. So they would take the bike in and it'd be dripping wet. And they would drip on the connectors and the connectors would get nicely corroded and break, and or they would kick the thing against the plumbing and somebody else would kick another one against the plumbing and you get a hum loop. So you get interesting faults where two machines could speak to each other, but not a third one. And that third one could speak to a different one, but not either of those

two, because of the hum loops on the wire. And made it look like there were collisions and machines didn't want to transmit to particular places. Quite a few times we sent the technicians just crawling along under everybody's desks, just fiddling the cables and making sure that they sort of work, but they got a bit of a headache for the admin staff eventually, because people kept buying these Suns and we kept having to add an extra subnets and extra bits of cable. And that got pulled through the offices. And so eventually the cable bundle must have been about that size [holding up hand to show roughly a 2-inch bundle] going through offices, and it got to a stage where we couldn't pull any more. We couldn't pull all the subnets through all the offices. So the admin people were constrained to having to put certain groups of people in certain parts of the building, because we could only network there simply because of the physical constraints. That was up to the late 80s and then round about '90-'91, we went for a twisted pair. And that solved that problem, we got a whole different set of problems from that. But that solved that problem.

Vassilis Galanos 20:20
It is extremely interesting; I think for researchers of today as well. From research I am exposed to, recent research, there is a tendency to observe that those who maintain the infrastructures get lesser recognition, that because of the immense amount of interface, we are less exposed to the technical limit of material infrastructure, and I wonder if you have any sort of historical reflections on this transition, the slow transition from you know, that being visible?

George Ross 21:01
I think even then people didn't realize just what was involved. I mean, the research people would just get a grant, buy another couple of Suns and expect that somehow they would be networked up. I don't think many of them ever gave very much thought to how it would be done. The people who were more on the hardware side of CS were aware of the problems. Some of them were very aware of the problems. But the further you went away from the hardware, bare metal people down to the more using it as a tool rather than using it as a thing to deal with, the further you got from the bare metal the less I think they were aware of what was involved in actually making this stuff work. And keeping it working.

Vassilis Galanos 22:05
Yeah, you mentioned that you as opposed to other academics, you were focusing more on the maintenance of the infrastructure. So were you closer to the metal side of things?

George Ross 22:18
Definitely yes. I would say certainly from networking side of things. I suppose I took over if you like with the killing off of the twisted pair. But somebody had to understand the ins and outs of the networking in more detail, and that ended up being me. I'm sure several of the other computing team at the time could probably have done it, but it ended up as me, and so I was kind of more involved at that point in the low levels. From then on really it was the low level, keeping things running at the low-level side. After I finished my PhD, it was round about then that the 68000 chips were coming out. 32 bit microprocessors. And a few people in the department had thought that's an interesting thing to play with, let's put together a modular system where we can just plug in bits and pieces. So various boards were designed, there was a half-Meg memory board and a processor board which had a 68000 on it. And another board which had our department ether on it, because Ethernet was new just about then, and the people who were doing the hardware, I was still not really involved at this point, people who were doing the hardware must have thought well that's a good idea, one single bit of cable just plug everything into it. So they did an Ethernet card for it. So this is essentially the Fred machine [aka APM],

and various people saw it and thought well that looks like a good thing. Let's jump on the bandwagon and not so much commercialize it but productize it perhaps. So it got turned into a definite product and eventually was used by pretty well all or I think most of the staff, and second- year upward students used them. And anyway that was quite a low level environment. Eventually there were various libraries, there was a graphics library done for it and things like that. So it gradually moved up a bit. But it was quite a low-level environment from my point of view. And that persisted until Sun Microsystems came along. That was an economy of scale thing. It's cheaper if Sun are building the things by the tens of 1000s. We're building them by the tens. So they can spread the cost across more machines. And they can make them cheaper and sell them cheaper. So once Sun Microsystems came along, I think at that point, probably the department drifted away from bare metal, unless you want to call VLSI bare metal, it drifted more away from bare metal at that point towards looking at systems rather than looking at machines.

George Ross 26:18
Yeah, that must have been 80s, late 80s, early 90s, somewhere around about that.

Vassilis Galanos 26:30
The time when the World Wide Web was made, maybe things were aligned. Very interesting.

George Ross 26:39
But the first thing I did for the department after I finished my PhD was rewrite the filestore systems. Because the Fred machines were coming along. The department had a central file server since 1976, it was based around an Interdata 70 with a couple of 60 meg disk drives on it. Like a washing machine size, disk pack about so big by so thick [demonstrating with hands, maybe 18 inches diameter and 3 inches high], half a dozen platters on it.

Vassilis Galanos 27:17
Like book volume.

George Ross 27:21
And so we had the Interdata system, but it was creaking a little bit at the seams under the load. And I think it was Hamish, Hamish [Dewar] and Rainer [Thonnes] probably thought, well yeah, we've got these Fred machines, these APMs, let's do a disk interface for that, and turn one of those into a filestore. So they did the disk interface, and the first job I got when I joined the department staff was to reimplement the file storage system. So it was everything, the protocol coming in on the wire, everything in the middle down to laying out the disk blocks on the disk. That was pretty bare metal really, and writing a primitive scheduler to schedule the processes that were dealing with all these things. And that actually ran for, it must have been running for about seven or eight years, those filestores. Eventually I think we had two 300 megabyte ones, one for staff and one for students. There was a 160 megabyte one which was kind of for development. The Winchester disk drives fit in a 19 inch rack, so they're 19 inch rack width, 19 inch rack depth, and they're probably six U high [https://en.wikipedia.org/wiki/Rack_unit], something like that. So that's giving you under 160 megabytes disk storage. I mean, now I've got a 64 Gig little XC card tucked into the side of the laptop there. And it’s smaller than a thumb. So that's how much storage has come on in the time. And we had a few other smaller ones as well. An even more development machine with a couple of 70 meg disk drives, five and three quarter inch size ones, they'd be what six inches by the time you put the cheek plates on to slot them into the rack. They fitted into the Fred machine boxes along with the processors and other cards

and things. That was quite close to metal, I suppose. As I say, we drifted away from the metal as people bought Suns, and it became more a case of keeping the network running and how do you manage the network of I can't remember how many Suns, we have the diagram, it probably has a few 100 on it. And how do you manage the network when you get to that stage? And as I say, the document I wrote goes through the various stages of how that developed from just thick net to thin net to, oh dear we can't carry on with this, let's go for twisted pair. So twisted pair to a patch panel and the technicians would patch the lines into hubs, depending on which subnet the machine was supposed to be on, which is a lot easier than people shunting between the offices. Because the techs could basically plug any machine into any subnet. And then going from that, because the patch areas got rather out of hand a few times, if you imagine people come and go. and the techs patch and unpatch. And eventually you have a rather a lot of knitting that's hard to deal with. So a couple of times, we said come in on Saturday, we'll pay you overtime, make a note of what's connected to what, rip the whole lot out and do it again neatly. And we had to do that a couple of times. Eventually, fortunately, again, it's economies of scale, HP ProCurves came out, modular switches which could be programmed remotely. So we could just patch everybody in once and for all, and then program the switch to say, right, this port is on this particular subnet, this other port next it's on a different one. And as people came and went, we just reprogrammed switches. So I wrote the code to do the programming of the switches. And that's been running for, I think it's still in use, that's been running for the best part of 30 years now.

Vassilis Galanos 32:26
Still the same code. Amazing. Yeah.

George Ross 32:31
I probably wouldn't do it quite the same way again, it kind of grew by accretion. But it does the job well enough. It means that the support staff can set people up without having to know too much detail about how the network works right down at the bottom level. You can just say, put that port onto that VLAN. Edit your file, close the Edit, change gets pushed to switch and it happens. And it's worked like that for, as I say 20 years at least, 30 years probably.

George Ross 33:07
But now EdLAN's getting redone and what they're wanting to do is make the network learn from what's plugged into it. So, if you're the network, you see a MAC address, and you think somebody's associated that MAC address with this particular machine, and it should be there [on this particular subnet]. Or this machine has given me this certificate, which says I'm so and so. So that certificate means that the configuration should be done like so. So the idea is to take the load off the people configuring the network, and let more of it happen automatically. Which I think makes sense most of the time. But I don't know, because it's about a year and a half since I had anything to do with this. But when I last looked at it there were still issues around things like the robot lab, for example. People tended to bring in robots and plug them in, and wanted to just be able to have a self-contained subnet they could just do what they liked with. Make a robot, plug it in, control it from another machine on the subnet. And that wasn't going to be quite so easy to do with this kind of network-that-learns situation. It would have to be programmed to say this bunch of ports are not self-learning, they're fixed to this particular subnet. And I don't know whether IS were expecting to be able to have to do this. And I'm not sure whether they eventually were set up fully to be able to cope with doing this kind of thing to any kind of extent. So even the new network has got little issues like that. But as I say, I don't know what's happening since. I noticed that Appleton Tower was getting its switches upgraded a few weeks ago, or some of Appleton Tower was. So I presume it must have moved on a little bit from when I left it, but I

think there are still going to be things that are going to be awkward cases, which people are going to have to find a good way to solve.

Vassilis Galanos 35:50
Interesting. Hopefully, this interview will help them as a hint?

George Ross 35:56
Well, I'm sure that the Systems Unit [part of the Informatics Computing Team] are quite well aware of all the problems. They were certainly aware of the issues when I left because we'd had various discussions and written various documents and discussion notes and things. So those I presume were still being used as background material for how people wanted to move things on. But that was a year and a half ago, and I presume it's all working, because nobody is phoning me in a panic saying it's not working, what should we do? So it must be, they must have worked out how to do enough of it, to keep it going. I don't know how much of this self-learning port stuff they've adopted. That was the eventual intention.

Vassilis Galanos 37:02
How have you been teaching during your...?

George Ross 37:08
Well back in the CS days. Before the formation of Informatics, as I said earlier, it was a much smaller organization. And there was much less distinction between academic staff, research staff, computing staff. We were actually not called computing staff, we were called “other related”, or something like that. It was a grade which was related to academic staff, it wasn't quite academic staff. So yes, one of the computing people [Fred King] actually did end up lecturing first year computer science, and being the course organizer for it. So there was a bit of crossover like that. I never did that, to that extent. Well, every post grad does tutoring, so I was doing tutoring as a post grad. But I carried on tutoring for a bit as a member of staff. I did things like invigilate exams, which computing staff wouldn't do now but they did do then. I remember invigilating I think it was written Chinese. You wouldn't be invigilating your own subject, you just go into the invigilating pool, and be scheduled to invigilate something. So I ended up in Adam House with I think about three or four different small exams going on. And I can't remember what the other ones were, but one of them was written Chinese. And I looked at the paper, and I could read the bit that said at the top University of Edinburgh, what the course was, Written Chinese. That was it. I couldn't read anything else. But because I was just randomly chosen to invigilate it. You just sit at the front.

Vassilis Galanos 39:12
So any type of courses, it was Chinese literature, or?

George Ross 39:16
I have no idea what it was.

George Ross 39:19
But as I say we were just given random things to invigilate. I did two or three invigilation sessions, and that was the one that stood out as the most unlikely one. The other thing we did for a while was for fourth year projects. We reckoned things were getting a bit complicated around the edges for them to see how things fitted together. So we started running a series of half a dozen lectures specifically for

the fourth years just at the start of their project. So I did I think it was a series of three on programming X Windows, for them. So there was one on how the applications fitted together, and a couple on toolkit programming for people who were wanting to design windowing tools as part of their projects. So we did that kind of thing for a few years. That was when we're still using Suns, so that would be 90s sometime.

George Ross 40:52
I'm sure I could look it up, I think I've still got the slides somewhere and they're probably dated. So they'd be in the loft. So I could, if you want to know I can dig them out of the loft.

George Ross 40:52
Yeah.

Vassilis Galanos 41:08
Just as an indication.

George Ross 41:11
This was when you were doing a slide you would either print it out and then put the slide through the photocopier with the paper printout, or else you would draw your diagram directly on the slide. I remember Gordon Brebner, he was doing the networking course, at one point, and I think he had a slide with about seven or eight different bits that folded in. So he had a basic slide, and then he added something, and then he added another layer, added another layer, and basically kept folding in transparencies on the top until it was about seven or eight layers thick, I think.

Vassilis Galanos 41:59
Probably the precursor to more contemporary slide transition mechanisms. But you have the same diagram kind of expanded.

George Ross 42:09
I mean, I never got to that stage. My transition was basically, you put the transparency down with a bit of paper on it. And as you talk to these people, move the bit paper further down to expose what are the slide.

Vassilis Galanos 42:21
What was the formal name of the device that was used to project the

George Ross 42:26
just an overhead projector.

Vassilis Galanos 42:30
I have some of those, but they're used for different types of presentation if you want to magnify something. And so my first university presentation was based on transparencies. Interesting, yeah.

George Ross 42:46
And the other thing we did, because we were a smaller department and you know, more unified, less differentiated perhaps, was that it was quite usual for the computing staff to propose and supervise fourth-year projects and MSc projects, and mark them. And so it was quite usual for us to supervise a

project, mark the project, go along to the board of examiners and talk about the student if necessary. Whereas after Informatics came along there was much more of, not so much us than them, but much more distinction between academic staff and research staff who did that kind of thing, and computing and admin stuff who didn't do that kind of thing any more. Which I think was a pity, because there was quite a lot that the computing staff could have brought to student projects, but it dropped out of circulation. And I had ideas, for example, I could have got somebody to do some kind of learning project about the network. Because we kept copious statistics about the network. Every five minutes, you grab every counter in sight from the network and stash it away somewhere. So traffic rates, error rates, link speeds, all that kind of thing. And it seemed to me that some kind of, it could have been an AI-ish type project, that somebody could have learned what the network did, and then monitor it, looking for things which were going wrong with it. There could have been quite a good project out of that but it never really happened. There could have been a machine learning project there, but because by this point we were much more disjoint from the project mechanisms that, I think I tossed it into the mix, but it never got anywhere. I think that's it. It's a pity that Informatics got so big because I reckoned it has think it fragmented into smaller groups again anyway. I mean, CS was small, but it was self contained. And everybody knew everybody within CS.

Vassilis Galanos 45:35
It's interesting, because people I've spoken to, from the AI side of the story, they talk about how everyone knew everybody, but not just within AI, generally, within the broader University of Edinburgh. I mean, maybe more in the Informatics kind of general domain, but, but also like connections with cognitive science and people like Christopher Longuet-Higgins, for example. And Donald Michie kind of became a centre of attention for and the main node of networking in the metaphorical sense, social sense.

George Ross 46:17
I suspect there probably was some of that going on with CS as well.

Vassilis Galanos 46:22
How did you experience that? Did you have the feeling that, you know, CS was more isolated?

George Ross 46:28
I wouldn't have said it was, I didn't feel isolated. Everything was small and self contained. I mean, before all the divisions and schools came along, you had lots of small departments. And so, as I said, I did statistics, the Statistics department was probably about six or seven lecturers and a secretary, under a Professor. So everything was that kind of scale. And then Statistics and Maths got merged, and probably something else. CS and AI and CogSci got merged, and Physics and Astronomy, and probably other things got merged, instead of about 40 different departments in the Faculty of Science. I mean, at one point, there was a Department of Physics and there was a separate Department of Mathematical Physics. I don't know whether it's just that a Professor would come along and form a new department, rather than slot into an existing unit,

Vassilis Galanos 47:43
postdocs and PhD students.

George Ross 47:46

And then you had your department with a Professor, four lecturers and half a dozen students. I think when I did Statistics we had one, maybe two PhD students. There were nine of us or maybe 10 of us on the MSc course. That kind of number. I remember when CS started up the MSc course I can remember the very first MSc student we had the first year we ran it, I can't remember what year it was, but we had one MSc student in the first year it was run. Now I suppose there are two or three hundred.

Vassilis Galanos 48:29
In your experience what separated, terminologically speaking, conceptually speaking, Computer Science from AI? And I'm also interested, you mentioned Statistics as part of doing Computer Science, also as an extension of Mathematics. So now many people say, well, machine learning and AI is just statistics. I think the same term is sometimes adopted by different. So, what was your experience around that? How did you feel at that time?

George Ross 49:15
I don't remember feeling particularly strange about the distinction between AI and CS and of course there was Machine Intelligence at that point as well. No, I don't remember it being particularly strange. I suppose CS had been coming up from the bare metal side of things that we were thinking about earlier.

Vassilis Galanos 49:48
Was that you think a historical kind

George Ross 49:50
I think it was this historical? Yes, it would. People had been building computers through the 60s and 70s, people were thinking, how will we use them. People like Sid Michaelson came in through the numerical analysis side of things. So he would have been approaching it from that side of how to use computers, starting off with mathematical and then drifting into the systems that made them work and operating systems, compilers, all that, networking. Networking I suppose was a little bit later. I mean, I was aware of AI existing, but not really aware of what they were doing, although some of my acquaintances, people who were doing CS round about the same time as I was, some of them were doing their fourth year projects in what you would call artificial intelligence, rather than computer science things. So possibly, because I'd never done the artificial intelligence courses, I was perhaps less aware of what there was in it from that point of view. I don't think I was ever thinking, oh this is different, it was just that everything was much more self contained.

Vassilis Galanos 50:14
So maybe it depended on just people, very random, but the different places people happen to be at different points in time. So in the 80s what many people talk about in these interviews is the Alvey Programme, the fifth generation computers, to computer science from Japan, and supposedly, according to most people, Alvey was the response to the Japanese computing [programme], some people associate that with AI as well. But did you have any?

George Ross 52:22
I remember the name, I didn't have any direct dealings? No.

Vassilis Galanos 52:34
do you think that the 80s was a good kind of period for Computer Science in Edinburgh?

George Ross 52:40
I think it seemed to be a good period. I'm just trying to think of how to phrase this, it was it was a good period for what we were doing in that in that kind of time period. The 80s would be microprocessors coming along. And because we had the department ethernet and the Fred machines and the file stores. And that worked quite well as an overall system for looking at the things that were around in the 80s. The teaching of computer science at that level, I think, it seemed to work well. The Fred machines did get developed quite well, eventually two or three different graphics card variants came for them for doing different sorts of graphics. And so I think they were an early networking workstation type set up, and then Suns came along. And everybody was buying Suns. So I think possibly at that point it became less distinctive, because other places would also be buying Suns. In the early to mid-ish 80s I think we were quite distinctive, because we had the microprocessor-based systems going and the VLSI design things as well. Once the graphics cards came along, you could display your VLSI designs on the graphics screens. So we had both of those distinctive things, I think, in the 80s

Vassilis Galanos 55:02
You mentioned a number of names that were there when you arrived. I'm interested in hearing you talking about the important people you've encountered Edinburgh. But also, if you have memories of lesser known people who have contributed in one way or another, but they haven't received enough recognition, maybe?

George Ross 55:32
Yeah. So I've already mentioned, the academic staff that were there when I came, it was Peter Schofield and Frank Stacey and Rosemary Candlin did the first year course. Hamish Dewar who died just recently. David Rees is still around, as far as I know. Rod and Robin. And then other people like Alex Wight and I'm trying to think who else was around at that point in the academic side of things.

George Ross 56:13
They were actually very good administrators as well as good lecturers. I think Sid didn't want to be head of department. So Peter Schofield took over as head of department. And he was a very good administrator, and he was very good at isolating the department from the vagaries of the rest of the university.

Vassilis Galanos 56:39
Okay, can you give an example of that?

George Ross 56:45
I'm trying to think of examples. But basically, we never really had to worry about things coming in from outside because PDS would field them and deal with them and sort them out as much as he could.

Vassilis Galanos 57:03
Was it as good so that you don't remember what the vagaries?

George Ross 57:06
Well, I don't remember what they were because we were all well protected. And then Roland came along, of course. Roland Ibbett. And Nigel was a PhD student that he brought along. And he was doing research into other kinds of networking. Who else was around at the time? I mentioned Gordon Brebner, of course. And Mark Jerrum, he went off to London. And I think afterwards, I'm not sure where

he is now. But we had people like Luca Cardelli, who came over and did his PhD with us, and then went back to the States. There were people like him when I was around.

Vassilis Galanos 57:58
So in terms of what made Edinburgh unique at that period, what kind of expertise did these people have, how did these kinds of modes of specialization interact with each other in order to bring about this kind of uniqueness? Why do you think it was a great place to study?

George Ross 58:25
I think it was a good place to study because we'd had the bottom up machine, bare metal side of things coming through, and a lot of people in the early 70s, 80s, early 80s would have been attracted by that. And the fact that EMAS was there. So the people who were teaching the course, had a lot of experience with actually writing the kind of systems that they were that they were teaching about. Then the VLSI people were there and that attracted VLSI type people. And then of course, the LFCS came along, and the more mathematical side of things, and that attracted quite a lot of people. I was going to mention another name earlier, I can't think now which one it was. But there were certainly people involved, not so much the teaching side, people like Fred King who designed the Fred machines and Rainer [Thonnes] who did a lot of the hardware designs, and even technicians, Peter Lindsay, who was designing boards for interfacing machines to things. He did the original link board for connecting these small computers together at a decent speed. There are people like them who were perhaps not so visible from the outside because they were concentrating on their particular area rather than doing the research and publishing it as well. Malcolm, of course, was my supervisor, Malcolm Atkinson. So it was 79. So he was with us for a while, and then went away, and then came back again.

George Ross 1:00:37
Just to remember who else was there, Lee Smith, I think was a name that I kind of remember, I don't remember where he ended up. Andrew Blake was with us for a while, I remember trying to get his network machines onto the network.

Vassilis Galanos 1:00:58
Interesting. We've interviewed Andrew on the basis of his AI contributions, but also CS. And Roland Ibbett as well.

George Ross 1:01:13
I remember going to a meeting. I can't remember when it would be, it must have been after we started putting in the Ethernet [for the Suns]. So there was a proposal that there would be a machine shared for teaching between CS and AI and Electrical Engineering. It was a Gould PowerNode. It was unusual for its time because it was a multiprocessor machine made up of less powerful processors, rather than what people tended to build at those times, which was a big, chunky, mainframe type thing. So I remember being at this meeting, there was a presentation and Sid was there and Roland was there, and I can't remember who else was there. And the guy had done the presentation and he'd been describing how the memory management worked so that the different processors could all work on the same memory system, how the cache coherency was controlled and that kind of thing. And I can't remember whether Roland looked at him or something, and he [the presenter] said, did you understand that? And then Roland said something like, well, I hope so because I invented it! Words to that effect. Actually, just an anecdote on Sid, Sid was one of these people who could quite easily give you the impression that he was not thinking about what you were talking about. We used to do fourth-year

presentations, when they were a small enough number of students that you could just schedule all the students to present for 20 minutes or something, and all the staff who wanted to could go along, and you could see Sid sort of looked as though he was nodding off, kind of falling asleep. And then come the end of the presentation, the one thing you could guarantee is that he would ask the most penetrating question of the student. He didn't look as though he was listening, but he was just taking it all in and concentrating really hard, although he didn't look like it.

Vassilis Galanos 1:03:45
There are a few people I've met who actually do that, and usually they're the most interesting contributors to their fields. If you have more anecdotes like that we're more than keen to hear.

George Ross 1:04:09
I think that was the most those two are the most memorable ones. You had to be very careful when you were presenting something when Sid was there. In fact, I would tend to warn my students when they were coming up to doing their presentations, I would warn them now Sid may look as though he's falling asleep but be careful because he isn't and he will ask you something.

Vassilis Galanos 1:04:42
Any other sort of maybe lesser known figures that you

George Ross 1:04:50
A lot of people kind of contributed in the background. I mean, the entire computing team would have contributed in one way or another, and there's probably about 30 or 40 names there that could be listed. I'm sure we've got the names on file somewhere. Actually, one of the things I think I still have is my office across in Appleton Tower because they haven't needed it for anything so far. And along the way, I ended up with a kind of repository for documents. And one of the things I acquired there was, I had a set of minutes from the Edinburgh Computer Users Committee, which I was on for a year or two. That must have been the 80s, thereabouts. I have a bunch of notes from what was then turned into the Gould Users Committee, which was the committee with CS, AI, Engineering. And the people from, what were they called at this point, I think they may have turned into University Computing Services by this point. Because they started out as ERCC, Edinburgh Regional Computing Centre, turned into University Computing Services, and then turned into Information System Services, or whatever that is.

Vassilis Galanos 1:06:31
was a gold power node?

George Ross 1:06:34
That was a Gould PowerNode. And so there was a couple of us from each of the departments and a couple of the people from the university computing services who actually ran the thing. Yes, it did get bought, and it sat in the CS machine halls. And as I think I said in the networking document, we connected to it directly because we were using a Gandalf multiplexer thing to connect all the terminals around the department. So we just connected the Gandalf directly to the Gould. AI had a statmux pair (statistical multiplex pair). And they had a terminal room in I think it was South Bridge. And they connected over a BT line to the statmux at one end and a statmux at the other, there was a kilostream line between them.

Vassilis Galanos 1:06:35

That was the building that went on fire.

George Ross 1:06:49
So I think it was somewhere in there. And they had their terminal room there and they connected to the Gould using this, it must have worked well enough. And EE connected to the Gould using, well, they bridged their ethernet to our ethernet, because somebody trundled their way through one of the heating ducts that ran between JCMB and whichever building it was in EE, and they pulled a fiber pair down this thing. So EE had, I think it must have been a bridge, a fiber bridge. So they had one end under our machine hall's floor connected to the Ethernet that the Gould was on. The other end was somewhere in Engineering. And they connected to the Gould over the Ethernet, over the fibres through the heating ducts. And that fiber actually ran for a long time, even after the Gould went away, we still kept that fiber running just as a backup link between CS and EE and the centre. So if one side went down, we could fail over to the other side. And that ran until about 2000 or thereabouts. And the only reason it's not still there is that they built a building on top of the thing, and that kind of broke the fiber. That was the Alrick Building. I think, one of the engineering buildings anyway got built and at that point the fiber got pulled out. But I mean, it served its time by then.

Vassilis Galanos 1:09:30
Do you think this connection, the technical connection, but also the social aspect of teaching together between AI, CS, do you think that this has implications for people communicating with each other as well? And sort of interdisciplinary?

George Ross 1:09:47
I wouldn't be surprised if it did. Yes. We certainly got closer together when facilities were shared like that. Certainly from the operational side of things, we certainly got to know the operational side AI people, the people over there better, as a result of having to use the shared facilities.

Vassilis Galanos 1:10:23
This is fascinating, because there are rather opposed views as to the impact of network technologies, communication. Some people say that because of too much online network, communication will become more isolated as humans, but in those early days, actually,

George Ross 1:10:43
I think it probably brought people together more.

Vassilis Galanos 1:10:47
Yeah, in different disciplines.

George Ross 1:10:50
I think, yes, because I remember, actually, we kind of got involved in meteorology as well, because they had a satellite dish up on the roof, which I think we ended up taking a feed from. But the Fred machines, APMs, whatever you want to call them, they got interested in those. And I think there was some coming and going between meteorology and CS, to use the Fred machines for meteorological purposes, for display and satellite downloads and things. Probably people just bumped into each other in the KB Union, or possibly even the staff club when the staff club still existed.

Vassilis Galanos 1:11:49

Okay, what was this Staff Club?

George Ross 1:11:51
I'm trying to remember how Chambers Street used to be laid out. If you come along from the Bridges end, you had Arnott's on the corner, and Al ended up there before it got burnt down. And some shops, then you had Adam House. Then you had the Staff Club, which was a kind of social place for members of staff to go. [Correction: Chambers Street House came next, then the Staff Club.] It was a physical place where they had quite a reasonable restaurant and a bar, lounge facilities and things like that. And if you were a member of staff, you could just join and hang around. You could go and have your pint with your pals, whatever. It probably worked better for the people in the centre of town. And I think people out at KB would just go to the KB Union instead. And then next to next to the Staff Club used to be Chambers Street Union, which was originally the Women's Union. And Teviot House was the men's union. And then when the Women's and Men's unions merged, Chambers Street House just became yet another union. They had the usual union facilities, a bar and a cafe type place. So when we were doing, Maths was either in Minto House or next door because there were various refurbishments going on. I think I was in both. I think I'd done math courses, or at least tutorials in both Minto House and the Maltings, which was next door. And we used to come out of, I think it must have been Pure Maths at one o'clock, perhaps it was a 12 to 1 lecture, head for Chambers Street House to get some lunch, and then head straight up to KB to get the two o'clock CS lectures. So we had a bit of time. Or else we'd come out of the lectures in Appleton Tower and sprint across the carpark that is now the Informatics Forum, because the minibus used to leave from the far side and if you didn't get there early enough, it would be full and you'd have to walk. Whereas if you got there early enough, you could cram in and get a minibus lift up to KB. And then we could have lunch up at KB those days in the KB Centre, which used to have quite a big restaurant in it, a big cafe.

Vassilis Galanos 1:11:53
It would be interesting to look up pictures for that period, I wonder if they exist? I never imagined the inexistence of the Informatics Building. It's so well entrenched to the landscape right now.

George Ross 1:15:08
Ah well when it was being built there were some cameras put up in Appleton Tower. Presumably you've seen the shots that we took from that? I can't remember who, I think it was the machine vision people who put up quite reasonable cameras, and the computing team had put up some not-so-good cameras, so we had six cameras pointing down at the carpark site when the Informatics Forum was being built. And so we were grabbing a shot every half hour, 20 minutes, something like that, for the entire building process. So you can see the whole thing, that's quite an interesting one. That's eventually what it looked like. [https://newbuildpics.inf.ed.ac.uk] So we've got a complete archive of all the shots, and then having an image every so often. So that's the car park. [https://newbuildpics.inf.ed.ac.uk/ATcamW/2005-05/06-16:00.jpg]

Vassilis Galanos 1:16:41
We will definitely include a link to that as part of your interview,

George Ross 1:16:44
And then you can see the ground works happening, the pile wall, digging out the basement, and then the various floors end up getting built and added in, and the cladding put on the outside, and at that

point things are starting to fall off Appleton Tower so they put some wire mesh over Appleton Tower to stop the windows falling off it.

Vassilis Galanos 1:16:53
2008, it's not too far away.

George Ross 1:17:19
Yeah, they started breaking the ground. it was still a car back in 2005. Round about November 2005 was when the piling machines moved in. And then the Bayes building, we put up a bunch of cameras for that Yes, we'd two cameras for the Bayes building and again, you can see it's going up, from when it was just a bit of ground with the paths across it and a bike shed in the corner. And that was 2016. And again, piling machines going in and then digging out a hole in the ground. We had one camera in 5.02, that must have been that one [displays a sample image]. And one camera, I can't remember whether it was in the Turing room or the secretary's office next to Turing room, but it was up there somewhere. So we'd want one at each side looking down at the site as the building was being built. [https://dtisite.inf.ed.ac.uk]

Vassilis Galanos 1:18:56
Fascinating perspective. I only arrived to Edinburgh two months after that, I arrived late August 2016. So it's just exactly what I missed.

George Ross 1:19:07
[This part of the interview involved discussion of some of the individual images from the archives.]
So by that time you'd be at that stage, they started putting the wall in the basement and they started excavating the basement, that's quite an interesting series of pictures because you find things that if you've looked at the drawings, you know they're there, but you never really see them on the ground. If you walk along the ground floor corridor at the East side of the building. So you've got the robot labs, and then wrap round to some offices and a fire door out. There's actually a service trench runs underneath that corridor, about so deep, with the various services that go from the basement round that trench and up into that must be Core C. And then they go up Core C to the various floors. So you find if you look at the drawings that things like that are being built. And in the basement, in the UPS room in the basement of the Forum, it was never really built a UPS room, but it was built as a kind of access to the end of that service duct. So you can see the end of the service duct up at the top right hand corner of that room, and loads of pipes and wires and things going into it.

Vassilis Galanos 1:21:10
Interesting. I want to return back a bit in terms of time period, I remember you said earlier when everyone got some Suns essentially, the significance of Edinburgh got a bit sidelined.

George Ross 1:21:40
I think maybe not sidelined, but I think the focus changed at that point. It moved away from bare metal hardware towards systems that run on existing machines that somebody else supplied.

Vassilis Galanos 1:21:58
And you also mentioned the transition from this more basic research to more product oriented.

George Ross 1:22:06

I think it's the same kind of drift away from bare metal to things which run at a higher level.

Vassilis Galanos 1:22:18
Do you consider that to be an obstacle for research, or something that it's just how things eventually ...

George Ross 1:22:24
I think it's just how things eventually went. Back at the beginning, people were kind of working out how to do the things on the bare metal. And nowadays, everybody just takes one or other version of Linux and runs it either a 386 equivalent or possibly an Arm box. And I mean, the tools are there, unless you're interested in the development of those systems, or how you would do things at that low level, most people I would suspect would just take the system and use it and not think too hard about how it runs underneath.

1:22:37
Somebody has to do the low level.

George Ross 1:23:07
Oh, yeah. Linux is pretty stable at that kind of level now, I think it's getting more things glued on to it, drivers and new hardware comes along. But I suspect the actual architecture is relatively stable now. People are not trying out different ways of doing things. With the Fred machines we had at least two different operating systems that ran on them, possibly even more, because there was the original one, and then there was the one that amounted to basically load a program and run it and speak to the file stores which was what most staff and students would use. And then we also had an experimental proper multi process system running on it as well, where you would run a process and send messages back and forward. The test version of the file stores, the hierarchical ones, made use of that. The original Fred machine filestore, I wrote a scheduler for that to do processes and message passing between the processes in a kind of basic way. So the filestore, the ether process, it grabs a packet coming in, a bunch of processes will interpret them and call the file system. And I think the disk driver had a separate process in it. And then there was another process or a bunch of processes that send stuff back to the client. So they were scheduled hierarchically, but not pre-emptively, just because it meant I didn't have to worry about consistency and concurrency issues. [https://history.dcs.ed.ac.uk/archive/os/APM-filestore/FS.1976/] But the revised new version, new as in late 80s version of the file store, we made use of the other multi-process machine and so they were proper real processes, you would send a message to the file system process saying get me a bit of file, and it would send a message back saying here it is. Or it would say, no, it's not, I don't have this file, or I don't have this. I can't tell you what this file is, but I can tell you who has it, go and send a message to this other guy instead. So it was it was using proper message passing between properly autonomous processes to do that. [https://history.dcs.ed.ac.uk/archive/os/APM-filestore/FS.hier/] So that was either the second or third version of the operating system, depending on how you want to call it on the Fred machines. But there was economies of scale, as I said earlier, people would port BSD, or Sun ported BSD to their hardware. LFCS had a VAX running 4.2 BSD which I managed for them for a while. But Sun ported BSD Unix to their Suns, and then they brought in System V or equivalent to the Suns. And after that, it was just commercial systems and you ran whatever you wanted on top of them. Though, I don't know whether possibly now the people making use of GPUs are possibly getting closer to the hardware again, or have been getting closer to the hardware again, for a while. Again, I suspect, they would probably have got the hardware and had do real hardware level programming, but systems have come along now, and they can gradually move up the tree again. As the things mature people are just using

systems as building blocks that exist, rather than as things they have to build themselves. There are people out there who are still interested in building systems. There's a very interesting mailing list called The Unix Historical Society, which is fairly active, and it's got a lot of UNIX old timers on it. Names that I recognize from back in the 70s, certainly 80s, still discussing how things worked and trying to, as we were saying earlier, get the history of the systems together. Because one of the one of the things we didn't do is we didn't really keep much in the way of historical change things. So we'd build a system and as I built the filestore, for example, I would just add a feature or change something and the old version would just get lost. Because we didn't have enough disk space to store that kind of stuff. So The Unix Historical Society people are trying to reconstruct old versions of various Unix systems, trying to figure out how things developed through the Unix timeline. I'm sure that I know there are things I would have, if I were looking back and had the space to keep them, I would have kept them now. But back in the 70s, 80s, we just didn't think about keeping old versions of things, tracking the changes.

Vassilis Galanos 1:29:20
It's something most people tend to do you know, when you're doing something like that. Sometimes when you're innovating, you don't realize, you just find a solution to a problem. You cannot even calculate its value, its historical, even financial value.

George Ross 1:29:41
From about the mid late 90s onwards, when we set up the CVS repositories and SVN repositories,

Vassilis Galanos 1:29:54
and these are exactly the things that we are interested in

George Ross 1:29:56
After that is, after that you've got the history. For example, that networking tool I built back in late 90s, whenever it was. I'm not sure how far back the RCS logs for that went. But once it got checked into CVS, there's a complete change history for that, right the way through to now.

Vassilis Galanos 1:30:21
Which, again, it's automatic, right?

George Ross 1:30:22
So it's automatic, so you don't have to think about it.

Vassilis Galanos 1:30:27
Which is a double edged sword, because you lack consciousness historically, sort of taking track, keeping track of the of the changes, and realizing what you're doing.

George Ross 1:30:38
I think that's one of the things that prompted the Edinburgh Computing History Project [https://history.dcs.ed.ac.uk] to come into being, I would say that was driven to a pretty large extent by Graham Toal, who was one of our students, and he would be two years or three years behind me. And he started collecting old versions of systems, going back to, in fact, we've now got some EMAS stuff on it. We have not very much VAX stuff, I don't think. We have a reasonable amount of filestore stuff on it. There is more filestore stuff that I have on CD backups. When the filestores were running, we would take a backup to tape every, I think it was every week or something, and they were stored in a

cupboard down in room one-six-something-or-other in JCMB, I can't remember the room number. And it was basically a rack with, say, fifty or sixty half inch tapes hanging off it. And we would have a bunch of tapes for the filestores, and a bunch of tapes for the VAX, and probably a bunch of tapes for the cstvax and possibly a bunch of tapes for the pyramid. I don't remember it having a tape drive maybe, maybe they did remote backups to cstvax on it. Anyway, there was a bunch of tapes in the tape store.

Vassilis Galanos 1:32:46
Is this archive still available?

George Ross 1:32:48
No, well, no, some of it is, and some of it isn't. The filestore stuff, John Butler grabbed a bunch of tapes before they all rotted and while we still had a tape drive that could read them and transferred them to CD. I have a copy of all of those CDs, I don't know who else has ever had copies. I think I may have one of the only copies left. But that, I think, is nice because the CDs are only 70 megabytes or something. So I just uploaded them on to the history site, in part of the site not exported to the web. And the intention was always that we would go through the CDs, and gradually move stuff that we could move from the private side of it onto the public side of it, but that never really went as far forward as it probably should have done due to mainly lack of time. And I would say it's only recently that Informatics has realized the value of its history. I mean, people would say, oh yeah, that's a good idea, but there was never any kind of real encouragement to do that, and never any time really set aside for it. So the CDs are still spinning on a disc somewhere, as well as sitting in my drawer at home. And I may have them on there [on the laptop I was using at the interview] actually as well. Can't remember. [There is. It's about 1GB, all in.]

Vassilis Galanos 1:34:40
That's good because my experience says that especially CDs produced after year 2000, they're not good quality either, especially those were made for mass kind of circulation.

George Ross 1:34:57
So there are things on those CD backups. Unfortunately, they're not complete backups, because it was a space thing again. We couldn't back up absolutely everything because we couldn't fit it all on the tape that we had, it would just take too many tapes. So there was a kind of a system set up where on the filestore, you could mark something for archive, and it would be backed up. Or you could mark something for vulnerable, it wouldn't be backed up. You would tend to back up source material. You'd not bother backing up object code, compiled things. So we have a reasonable amount of source material for the Fred machine, systems that ran on the Fred machines, not all of which has been transferred to the public side yet. Because, partly because I didn't ever have time to do it, and partly it's a case of untangling all the personal stuff from the departmental stuff. Because, although it was a kind of staff filestore, there were directories on the staff filestore which were for specific things, like the VLSI software or the ESDL board design software, that kind of thing. But scattered amongst all that lot there were directories for people, and it would need somebody who knew who was who or had a reasonable recollection of if you see a username, this is likely to be a person, or this is likely to be a non-person, an agent, an actual package user, if you like. So we'd need somebody who knew a bit about that, and who could then work their way through the various different backups, because we've got backups over three or four different years for each of the different filestores. And so we need somebody who understood how the backups were made, and then go through the backups and be able to pull forward the relevant things into the history public website. And also, there is in existence a backup of ECSVAX, which

Alastair Scobie has access to. I don't think he's ever had time to do anything with it. And again, there's the same issues of pulling out personal stuff from private stuff. So if the School really wanted to get more of its history online, in the way of systems and sources of the systems, then it needs to commit somebody's time, and somebody who's aware of the data protection issues. There'll be privacy issues, and there are probably possibly proprietary issues as well. I don't know whether all of the stuff on there was ever meant to be public, as it were. Because I remember there was a slight debate about when Hamish wrote his Imp80 compiler, because by this point, he had split out Clan Systems with Igor Hansen, and the two of them were doing laser printer development, amongst other things. And I think I have a vague recollection that Hamish wrote the compiler and brought it in on, there was some kind of condition that we could have the compiler, but not the source of the compiler, or not all of the source of the compiler, anyway, because it was proprietary. And there may be other things in amongst the archives that we have of the filestores, and what's left of the ECSVAX archive, I don't know how much of it there is. I did get Alastair to pull some stuff off it for me two or three years ago. So I did manage to retrieve my TCP/IP code for running on VMS. [https://history.dcs.ed.ac.uk/archive/os/gdmr-tcp-ip/] I already had the Fred machine version, because it was code that would run on both with a bit of conditional compilation. Backtracking a little, we had various islands of connectivity, if you like, at one point. So we had all the Suns and BSD boxes, and they all spoke TCP/IP as did the Gould PowerNode. We had ECSVAX and Gordon and Paul McClellan had written X.25 based software for that to connect to the wide area network, so we could go out to the wide area network and back in again to cstvax, because it had a connection out to the wide area network as well. But it wasn't easy to get directly between them. And the Fred machines were completely separate. And we also had a bunch of machines in the machine hall for ERCC's commercial arm, and they were developing Coloured Book software. And that was completely different again. So I wrote a TCP/IP stack, not a totally complete one, but a sufficient one that would speak to the BSD boxes. So the Sun's and the 4.2 BSDs and the pyramid and other things. And it did TCP, it did UDP, it knew about subnetting, and it knew about basic routing, and it could do bi-directional rlogin, rcp, bi-directional printing. So you could sit on your Sun workstation, do lpr blah, and they would throw the thing at ECSVAX, which would spool it onto the VMS spooling system, which would then throw it out to a printer was connected to ECSVAX, or you could do it the other way around. You could be sitting on your terminal connected to ECSVAX through the Gandalf multiplexer, and you could say print slash-whatever-printer, I can't remember the qualifiers and modifiers now, but you could basically specify a printer which was attached to a Unix box somewhere, and the print spooler on the ECSVAX would know to throw it to the print spooler, through my software, through to the print spooler on the Sun or BSD box or whatever, and it would print out on that. How did we get to that? Islands of connectivity? Oh, yes, I remember. So, yes, so the code that I wrote for ECSVAX and the Fred machines, I always had the Fred machine copy of that, because I got it from the filestore backups, but I'd never had the ECSVAX VMS copy of that, and I wanted a copy of that because, when I wrote that code, I was originally using VMS mailboxes, which was built into the system, so you set up a mailbox between two processes, and you send a message into the mailbox, and you read the message out to the mailbox, and I realized that it was a bit slow, and it was lacking some facilities. It was basically copying buffers about too many times. So I was looking for my version because I'd written a pseudo-device driver for VMS which went a lot faster, because it got rid of most of the copies. Instead of copying from buffer to buffer to buffer, the driver would say, oh, here's a buffer, I'm just going to hand it to this other process, and I'll do the accounting for the two process information blocks to make it look as though the buffer belongs to there and the quotas get adjusted. And the driver would pass buffers around and just adjust the process quotas behind the scenes to make it all work. And that went a lot quicker, but I'd never had a copy of that. And then I managed to persuade Alastair to give me a copy of my files, off the file store, off the ECSVAX backup. And fortunately it included the

TCP/IP stack for VMS as well. So that's also now on the History site, but there'll be other things on that ECSVAX backup, and somebody will have to persuade Alastair before he retires to hand it on to somebody else. And for the filestore backups, somebody will need to find a way of getting an appropriate person to do the transfer of the remaining stuff and the separation.

Vassilis Galanos 1:44:55
I wonder if there is a way to bypass the legality of open code versus closed code through.

George Ross 1:45:06
It might be you could get whoever the Head of School is these days to say, well, if it's on the backup it belongs to the School, and therefore, I could put it onto the…

Vassilis Galanos 1:45:17
or historical preservation if you managed to use an argument based on cultural heritage. You know, Edinburgh history is protected by UNESCO. So yeah, there are ways to

George Ross 1:45:32
I mean one of the reasons I've still got an Inf laptop and still got a visitor account is that I got possibly put down as a consultant when I retired, and they weren't sure whether they would need people to be consulting with me to see how the network worked. In fact, that never happened, and nobody has phoned me up. I still have my visitor account, and if somebody wanted to pay me for a day a week or something for a short period of time, then I'm sure it could be arranged.

Vassilis Galanos 1:46:12
I mean, that this is a possibility. I'm speaking of which I'm also aware of time passing. Although this conversation is just amazing, I'm just tempted to continue asking you things, but this is the right point, I think, to move to your presentation of the archive.

George Ross 1:46:37
I can guide you through what we have.

Vassilis Galanos 1:46:41
I would say, before we do that, maybe, would you have any closing remarks, any advice to give to new computer scientists in the new century, something you carry out from being at Edinburgh University for these decades. You know what kind of inspiration?

George Ross 1:47:03
I mean given that I was there for 40 years, I would say I was not doing the same job for 40 years, it was changing constantly. And I would say that if somebody's coming along now starting out in computer science or AI or whatever, be flexible, because in five years time it'll be different, in ten years time it'll be totally different. And in 40 years time, you'll be looking back thinking, oh, yeah, we did it that way. So I think it's going to change, keep changing, and be flexible, is what I would say to people, versatility is where it's at. But also, don't forget what was done in the past, because you can still learn from that.