An interview with
IAIN S. DUFF

Conducted by Thomas Haigh
On
30 & 31st August, 2005
In East Hagbourne and Chilton, United Kingdom

Interview conducted by the Society for Industrial and Applied Mathematics, as part of grant # DE-FG02-01ER25547 awarded by the US Department of Energy.

ABSTRACT

Iain S. Duff explores the whole of his career to date. Duff grew up in Glasgow, attending The High School before earning a first class degree in Mathematics and Natural Philosophy from the University of Glasgow in 1969. He discusses the mathematical curriculum there, and recounts his early computing experiences as a summer fellow with IBM. Duff won a Carnegie Fellowship for postgraduate work at Oxford, culminating in a D.Phil. conferred in 1972. At Oxford he worked with the numerical analysis group and wrote his thesis on the analysis of sparse systems under the direction of Leslie Fox. While working toward his degree, Duff was already collaborating with Alan Curtis, Michael Powell, and John Reid at the nearby AERE Harwell nuclear power research center. His interest in sparse matrices was cemented by a 1971 visit to IBM's Yorktown Heights laboratory for a conference, where he first met colleagues such as Alan George, Ralph Willoughby and Fred Gustavson. In 1972, Duff went to the United States on a Harkness postdoctoral fellowship, spending time at SUNY Stony Brook with Reggie Tewarson and Stanford, where he got to know Gene Golub. In 1973 he became a computing lecturer at Newcastle University, before returning to Harwell as a researcher in 1975. He continued to work on sparse matrix methods, establishing himself as a leading researcher in the area, through pioneering work with John Reid on multifrontal methods, a series of widely read review articles, and an influential book written in collaboration with Reid and Al Erisman. Duff rose within Harwell, becoming leader of its numerical analysis group in 1986. He explores in detail the development and composition of this group, paying particular attention to the Harwell Subroutine Library, a collection of numerical subroutines developed within Harwell from the mid-1960s on and widely distributed to other sites. Duff explains its history, evolution, areas of strength, sources of routines, documentation methods, distribution practices, relationship to the NAG library, and attempts at commercial exploitation. In 1990 Duff led his group from Harwell to the nearby Rutherford Appleton Laboratory, as government cuts and shifts in priorities threatened its future at Harwell. Since 1987 Duff has also led the Parallel Algorithms group at CERFACS in Toulouse, and discusses the origins of CERFACS, the composition and functioning of his group, and its involvement in the MUMPS and PARASOL projects to create parallel multifrontal solvers for sparse systems. Duff worked on various forms of BLAS, including contributions to the Level 3 BLAS specifications (with Jack Dongarra and others) and development of the Sparse BLAS standard (with Mike Heroux and Roldan Pozo). Duff also reviews his roles in a number of societies, including SIAM (as a Council member and Trustee), IMA (including editorship of its Journal of Numerical Analysis), LMS, ACM SIGNUM and SMAI, and his work as a consultant on supercomputing issues to the governments of countries including Malaysia, Sweden, Germany and the USA.

HAIGH: Thank you very much for agreeing to take part in the interview. I wonder if you could begin just by talking in general terms about your early life and family background.

DUFF: Yes, sure. Well, thank you for inviting me to participate in this. I don't know, maybe it could be regarded as an honor, but it's certainly very interesting to talk about one's life. I'm very impressed by the background that you've already provided. You seem to know more about me than I do about myself.

I guess I'm very proud of coming from the city of Glasgow in Scotland where I was born and brought up, went to high school and took my first degree at university. The family background, my father was in life insurance and his father was a shipbuilder or a tool maker, in fact, with John Brown's Shipyard in Glasgow, so our roots are pretty strong in the classical tradition. From an early stage, I found mathematics easy, I would say. I'm not sure how excited I was; it was just the subject that you didn't have to do any homework in and you could still get very high percentages and good marks, which made it attractive to me so I could do other things like play sports and be an ordinary kid when I was at school.

The early education was in local primary schools where we lived in Glasgow. Then I went to what's called the High School of Glasgow, which was an old-fashioned grammar school, which is to say it was selective in terms of people having to pass entrance exams to go to it. But it was a corporation school at the time. It has since, in fact, gone private, but it was a regular state school at the time I was there. That school, of course, emphasized quite a lot in the classics, as grammar schools would do. So by the time I came to what they call the Scot Bac now, or the Scottish Higher Certificate Exams, which you sit in five or six subjects, rather than the more specialized A Levels in England, I was really left with the choice of doing either Latin and Greek at university, which is what I was doing in the highers, or doing math, physics, and chemistry, which were the science subjects of choice at the time. Again, the choice was actually not an easy one, which may be a bad thing. Maybe it shows that I'm not just a total natural mathematician. But I opted for the science subjects and mathematics for the same reason, that I found them easy and you didn't have to do a huge amount of reading or preparation for the exams.

HAIGH: Would that have been a decision that you made when you were 15?

DUFF: No, this would be a little later. 15 would be the kind of equivalent of O Levels. In the Scottish System, the decision is delayed because you do quite a lot more subjects at the higher level of school, so this is a decision that I would have made at 17 or 18, really, about a year before I went to university. I stayed on for the extra year at school because I was enjoying it so much and I was running a lot of societies, involved in quite a lot of activities at the school, although I could have left for university a year early.

So as I said, it wasn't exactly a toss-up between the two sides of things, but it was certainly something I looked at really carefully. And I still like things like history and classics, almost as a hobby.

HAIGH: So while we're talking about these topics, is there anything you would say from your experience and upbringing in Glasgow that had a particular effect on the way that you approached your later career, that you think had shaped your personal philosophy?

DUFF: Well, I think people who are sort of energetic and honest and upfront, and I guess even friendly and enthusiastic, is an important thing in the sort of environment I was brought up in in Glasgow. I would like to think that that has continued and influenced my whole life since then. I do get very enthusiastic about things I do—about mathematics, about new maths analysis, about research topics that I'm doing nowadays. That's something that I've always been enthusiastic about. I was enthusiastic about mathematics in school, for example. So this kind of general flavor of life was very strong there, I would say. I was going to say in the beginning, that I value very much the fact and the background of being brought up in what to my mind is still one of the best cities in the world! A bit biased, but that's right.

And of course, it was automatic to go to the University in Glasgow. Oxford and Cambridge just don't really feature in the Scottish psyche. At that time, at least, anyway. So the best thing that you could do was to go to your local university and, in fact, I stayed at home during that time, which might sound unusual. But in fact, one of the benefits of being at home is that the university year is then a 52-week year, so all activities continue throughout the summer, unlike dormitory or residential or campus universities where they are only functioning for the term time. So there were good and bad points about staying home, of course, but I think there were a lot of good points. Well, there were bad points; it wasn't until I went to Oxford that I had any sort of independence in terms of living separately.

HAIGH: You had mentioned that one of the things that you had liked about mathematics was that if you were good at it, you didn't have to do that much work. Did you also enjoy those technical kinds of hobbies like chemistry sets or building radios?

DUFF: [Laughs] Well, I tried to build radios; I wasn't very good at that. But chemistry sets, yes, I did, until I nearly blew my eye out one time by mixing two chemicals that shouldn't have been mixed in the test tube. Fortunately I was already wearing glasses. I had glasses from the age of about 15 with short sight, and that probably saved my eye! So yes, I did like that type of thing. You know, the degree at Glasgow is called "Maths and Natural Philosophy". Natural philosophy to most people just means physics. I did actually like the physical side more than the chemical really. I did the chemistry in the first year at university and biochemistry in second year, but the maths and physics or natural philosophy I continued. That was interesting, and the combination was interesting, how the mathematics could be used. The problem really in that course, was that physics always wanted more mathematics than you had actually done in mathematics and the maths wasn't taught very well by the physicists, so it was a bit tricky. But when you realize what the maths was, later, suddenly the physics became more clear. Things like wave theory and Fourier analysis, that type of thing, were very much coming into physics before they were taught in math.

HAIGH: All right. So just before we move on to talk about your undergraduate experiences in more detail, I think I have a few more questions about the High School. You said that you were enthusiastically involved in running a number of societies. What kinds of things were those societies devoted to?

DUFF: Well, some of them were a bit geeky like the chess club. We had one of the best chess teams in Scotland at the time. I think my peak was coming in third in the international Scottish championship for people under the age of 18. So I was very keen on that. I haven't played chess more recently because it's just got too competitive. And what's called the "literary society" which was really the debating society in the school. I was, I can't remember, secretary or chairman, but I was effectively running that for the last year at school. We had debates with other schools and we entered into competitions with this type of thing. The school to this day still has one of the best debating clubs in Britain and wins prizes. I wouldn't say I was the best of the debaters, but I was sort of running the show and we had some very good people who entered into the competitions. That continued, of course, into Glasgow University. The debating chamber is probably one of the best in the UK for political debates, but we'll come to that shortly.

Other things, I guess I was involved a bit in cricket—that was the main sport. Rugby I was very keen on when I was young, but having glasses in the days before contact lenses were really a good option kind of put an end to the rugby career. So I played rugby until 15 or so. Well, in fact, I played a little afterwards without my glasses, but again, it got so bad that I wasn't seeing the ball. So cricket I played a little bit of, but never got to the extent of playing First XI really. I guess those were the two main sporting activities at that time

HAIGH: Would this school have been considered one of the most academically challenging in Scotland?

DUFF: Yes, well when we were applying the year before going to Glasgow University, there was a thing called The Glasgow University Bursary Competition. I was competing in that with other people when I was in sixth form. Of the first ten places in the Bursary (and this is open to anybody in the UK, but it's mostly Scotland to be honest), five went to our school. We had eight in the top twenty or something. And the McCrie Fellowship was a prize. I was somewhere in the top ten. The first place went to a pupil at my school, Philip Chalmers, and he was more in the sort of classics side. We certainly cleaned up that bursary that year, so I would say we were certainly the top school in Scotland at that time. You can't really evaluate with the rest of the UK because the educational system is different in England. And, as I said, we also won one of the chess tournaments when I was in charge of the chess team, so we were quite good there as well.

HAIGH: How did the Scottish undergraduate degree system work at that point? Was it that you started with a larger number of subjects in the early years and then you would finish with just one in the final year?

DUFF: No, the scheme is very similar to this day. And that's one thing that England has thought about moving towards. Let's suppose you're somebody that's heading to university. Then you would almost always be doing at least five separate subjects at what's called the higher level, which is the university entrance level. I actually did eight, which is rather unusual, to be honest. You tend to be sitting that a bit younger than in England, about 17 to 18 years old. Then you could go straight up to university. The reason for that is that you actually start the secondary year a year early. So I've forgotten now, it's either 11 rather than 12 or 10 rather than 11. It's one year earlier than in England. So the whole thing is much more general. It goes all the way through.

In the sixth form, which is an optional sixth form as I said, because you can leave for university at the end of the fifth form, aged about 17 to 18. The sixth form, one can specialize more, and I did three maths subjects. This was probably the beginning of the deciding to do math. It was dynamics, elementary analysis...

HAIGH: And geometry.

DUFF: Well, okay. Just amazing! Well done! Euclidean Geometry, that would have been of course, then. Yes, at that time, I was getting more into math, and I guess I had probably decided to do mathematics by that time. So you do specialize a bit more in the sixth form. Even in the sixth form I was still doing Latin. It was kind of compulsory in a grammar school in these days.

It was elitist, one must say. Because university education doesn't mean anything nowadays, although that's my bias. It was elitist based on ability. But then to go to university you had to have the middle school (that's 15 year-old, O Level/GCSE) pass in a modern language (that caused me a little trouble because I was only doing French one period a week), have a pass in mathematics in that level, irrespective of what you're doing, and I'm pretty sure that a pass in Latin was still obligatory at that time. This is at the lower level, not the higher level. And English, of course, you had to have. So there were a lot of requirements for university entrance for whatever subject you did, which of course, absolutely don't exist anymore. I think that gave people that were quite rounded in education. Some people like, for example, my wife, who went to university at the same time, had terrible trouble getting her GCSE in mathematics, and it wasn't uncommon that that was the hard subject. But you had to have that to go to university back in 1969.

HAIGH: Was it the case in Scotland, as in England during the post-War decades, that the university fees and some living expenses would be covered by the local authority grant?

DUFF: Yes, it was altogether much more generous than nowadays. There was never any question, at least if you were a British citizen, of paying any fees. In fact, I think even overseas students were free at the time I went to university, and then they brought in fees, I remember, protesting about that during the time. That would be in '65 and '69; they first brought in fees for overseas students, which caused a big uproar in the student body.

HAIGH: I had asked about how the university curriculum worked. So you graduated in "mathematics and natural philosophy." Was it mathematics and natural philosophy all the way through your whole time there?

DUFF: Yes, well, I did maths, physics, and chemistry the first year, and maths, physics (physics and natural philosophy are the same thing), and biochemistry in the second year. I guess it was just maths and physics for the third and fourth years. It's a four-year course, the honors course in Scotland. So I did maths and physics just for the last two years, I believe. I don't think I did anything else in the third year.

HAIGH: So would everybody enrolled in that degree program do the same things?

DUFF: No.

HAIGH: Or was it like the tripos system where you choose a different thing every year?

DUFF: Yes, the other topics, chemistry and biochemistry in my case, were completely optional. You could do anything that the science faculty would allow you to do. So you could do psychology even, which was included in science, or geology. Lots of people did geology; it was quite popular.

HAIGH: So can you talk about the kinds of things that you would have covered during the degree, with particular reference to applied mathematics and numerical analysis?

DUFF: Well, right, that's the thing. There was no computer. Zilch. Completely none. I don't think there was even a computer at the university at that time. So that wasn't on the agenda. Numerical analysis, which may perhaps by that time have been taught at some places in the UK, that wasn't taught at all. We may have done something like Simpson's rule for evaluating an integral—that's a numerical method. But that was more involved with like, series expansions or showing how you can converge two limits and this kind of stuff. It was done in that vein rather than trying to compute an integral.

Ian Sneddon was one of the professors, who unfortunately died about four or five years ago. But was actually one of the most well-known, international known professors in applied maths. He taught us Fourier analysis and that type of thing, which was actually extremely important for the physics (but unfortunately lagged behind the time that you'd need it for the physics). But the strength of the University of Glasgow Maths Department at the time was really in analysis, primarily the department head Robert Rankin. Alas he died quite recently. This was all the Kappa/epsilon stuff and that, and he had just finished his huge, big book on it which was from his lecture notes from previous classes when we arrived, so we had that to study. That was really the critical thing for people doing maths. So in the second year, you have to decide what honors course you're doing—it's a four year course, so the shoice is made at the end of second year. All along I was thinking of doing maths and physics; it wasn't an issue. But you had to do well in the exam in this course, which was the first serious maths exam, I would say, one did in one's life. I would say that of people doing maths, only less than a quarter—it might have been 20% or 15%—got more than 50% in that exam. For people who were used to getting 100% in maths exams, that was sobering. I think I got mid-70s or 80, and I think I was about fourth in the class for that exam. I was quite pleased and, of course, admitted to the honors program after. But that was the make-or-break point in peoples' maths career. And there was no sympathy. If you didn't get it, you were out on your ear and did some other subject. So that was the point towards the end of the second year in the maths at Glasgow that that was decided. And that was Robert Rankin, who was head of the department and the main professor there.

There was also a fair bit of linear algebra, but not numerical algebra, in which I work in, but vector spaces and the like. That was run by Dr Graham. I think he was just a doctor, because we didn't have many professors in those days. Dr Gillespie also did some of that. So these were the mainstays of the maths department at the time. They were very enthusiastic lecturers, I must say, which helped a lot.

The physics department, they were also very energetic and quite involved in university politics. Professor Ivor Dee, who was very keen on bubble chamber work, and Professor Gunn, and Bishop, who I met many years later at the European lab in Ispra in Northern Italy. He was either head of the lab or head of research. I happened just to meet him by chance. He claimed to remember me, but I think he had done some homework, knowing I was there.

So yes, the atmospheres were quite nice. There was certainly them and us. The students weren't allowed to use the elevators for the staff and that kind of the stuff. This was in the old days when students and staff were different. But having said that, the atmosphere was quite fine and we really didn't think anything of the discrimination. That was just natural then. As I said, the real problem work-wise was that the physics ran ahead of the math, and that was very frustrating, especially when you're taught the maths badly by the physicists.

HAIGH: Did you have any relationships with professors or with other students which were important later on, other than those that you have mentioned?

DUFF: I guess I had a good relationship with most of the professors. I wasn't hesitant about coming to see them offline to understand stuff in the lectures, and they were always helpful people, like I mentioned Drs Graham, Gillespie, and McHarg, I think it was, some woman who did functional analysis and stuff, and I guess that was important a little bit for the numerical side because functional analysis is one of the areas I did in the M.Sc. later on. So no special relationship, but just generally, it was very good.

Students tended to be very involved in politics and social work—I would say social work, but working with "Free Nelson Mandela" kind of stuff and that sort of thing, which was very popular at that time. So I was very involved in that, and I think I was unique in being on the committee of both the conservative party and the labor party in the Glasgow University debate team, moving from conservative to labor as I became more committed to social issues. I'm still extremely strong in politics at the moment. So that was very important. I was also on the student representative council. I was elected to that in a very bitterly contested election, which was partly political and partly just to be involved in the student body. I played a lot of bridge when I was at university. I gave up chess for bridge, essentially. But really social bridge.

HAIGH: So was it unusual for science students to be involved in student government and such?

DUFF: Yes, very. Don't forget, I'm talking about the period over 1968, which was when the student revolutions were happening all over Europe, and that was a very exciting time to be a student. I also visited America in 1971, which was still during the Vietnam War. The Vietnam War was a terrible time for America, but it was actually a very exciting time to be a student, because there was really an issue that united students in a way that almost nothing has since. So in the middle of all this, in about 1968 or '69, we were inviting Daniel Cohen-Bendit to Glasgow and address the student union. I said I was on the representative council inviting them, and in fact was in the subcommittee organizing it. It was called "The Teach-In on Higher Education". I had forgotten about that until you brought this up. That was completely fascinating, contacting all those people that are still famous nowadays—Daniel the Red and people. At one point, I got called to account by Professor Gunn, the physics professor, who more or less lambasted me and said, "It's okay for you art students to have all this time to do things." I had a little debate. I think I did confess that I was one of his students in physics, and he was quite astounded by that. So it was very unusual for a science person to be as involved as I was. I had some very close friends in Glasgow who had actually come up from the High School with me. I mean, of our class, like 30 people in the top class at high school, which was the size of the top class, I think every single one went to Glasgow University, which was the natural thing to do. So I still had all these contacts, and of course, many of them were in the arts topics.

HAIGH: So was there a strong rivalry between Glasgow and Edinburgh?

DUFF: [Laughs] Yes, of course! I mean, we just discounted Edinburgh when we were all there. There was a strong rivalry, ironically enough, between Glasgow University, and what we called at that time Strath Tech. Strathclyde became a university just a little before I went up. It was in 1963 or ’64 I think it was given university status from being otherwise Scotland’s top technical college. We used to have great battles with them at rag days. Certainly at that point we looked down upon them as not being as good as Glasgow. This has clearly changed in recent years, as well as the fact that I’m a visiting professor at Strathclyde now. Edinburgh, yes, Edinburgh wasn’t dangerous. We didn’t have the same direct pitch battles in the streets like we might have had with Strathclyde, but we certainly didn’t have so much time for them. Because of rugby internationals, I used to get over to Edinburgh regularly, and if we ended up going we’d have a dance in the Edinburgh University union. We had a reciprocal arrangement with them. I did have an Edinburgh girlfriend for a while. So we did have some contact. But certainly there was a big rivalry between the two cities and there still is to this day. Glasgow is the main city in Scotland, but Edinburgh is nominally the capital and, of course, has used that political clout to get a lot of things it shouldn’t have over the years, including the new Scottish Parliament, which is rather famously overspent, which was granted to Edinburgh, but might have been better for the budget to be in Glasgow. So there’s a lot of rivalry as between the two.

HAIGH: So it appears from your résumé that your first exposure to computers would have been during summer work with IBM while you were in undergraduate. Is that true?

DUFF: Yes. So what happened was this IBM Fellowship came about from some advert somewhere, and it looked very attractive. You asked about how one was supported. In fact, I got more money than most students—slightly more, not a lot—more money because I got one of these IBM Fellowships, and I think I was the only non- Oxford and Cambridge graduate to get the fellowship in the whole of the UK. It was UK-wide and they had seven or eight, maybe, fellowships in the UK each year for going as an undergraduate to university. At least in the year I was there, I was the only non-Oxbridge person to get it. Part of the deal, I don’t know if it was completely compulsory but I was very happy to go with that, was to work for IBM during the summer vacations. So I worked in various places. The memorable places were Toronto—that was my first visit to North America, down at the IBM Data Center on King Street East. I worked in the Glasgow Data Center and I worked in Birmingham, again in data centers. So I wasn’t doing much heavy numerical computing, although because I was doing a maths degree and because I was interested, they quite often gave me little jobs to do which were nothing more than evaluating functions or something. But my first main experience was in the Glasgow Data Center with IBM 1401 equipment. The first programming language that I ever did was 1401 Autocoder, which is a step up from machine-level language, but was kind of primitive by today’s standard. I also did RPG, which apparently could get me some money today because they’re still looking for people that can do that program generator. But again, it was a silly language. It was a spreadsheet language. It was almost like Excel on cards. That was the languages I started in.

HAIGH: How did you react to this experience in computing? Was it something that you enjoyed a lot?

DUFF: Yes, I did. I find it quite exciting to be with computers. It was great on a 1401 because you used to be able to look at the registers and single-cycle it. You could actually see by lights what was in these registers. It would show on the display. You could press a button, single-cycle it, look again and see what happened and I worked through the computation that way. Well, this wasn’t heavy numerical analysis; it was mostly data processing. But it was really just amazing, I thought. I still am amazed by computers and what they do. That was my first exposure and it was really quite exciting. I really enjoyed working with them. The people in that area at the time, and this was long before management consultants, and all these other guys got involved and ruined the whole topic—these was really people who were excited about computers and it as a fairly new thing in business use. The atmosphere was really quite exciting. IBM were really a bit condescending to the customers. I was definitely an IBM person—there was really no question. I had to wear a suit and a white shirt and a tie, even though I was a graduate student. And we more or less were told (I don't think this was in the IBM manuals), that all of your customers were idiots and clowns and you have to be really nice to them and that, because they're intellectually inferior. That was the IBM ethos. I don't know if they've still got that. But they did recruit well, and the IBM people were good. So it was great working with them, working with people who really knew what they were about, and not just pretending to know what they were about. These were real workers. So it was good.

HAIGH: So the data center was running as a service bureau, was it?

DUFF: Yes, essentially that's right. I used to go out to customers and help a little bit. When I was in Birmingham I helped somebody who had some mechanical things, and it was quite complicated. I really can't remember exactly, but it was geometry, essentially, analytical geometry that was involved. So that was the closest thing that I got to reasonable maths in it, but mostly it was payrolls and other things of that kind.

HAIGH: And within the data center, would there be separate operators different from the programmers, or would programmers be able to run their own jobs?

DUFF: Well, the elite in the data centers at that time were called systems analysts, who were sort of super-programmers, so they had gone through that and were kind of half-managers and half-programmers, and I was associated with them. I was working with the systems analysts, but I was doing a lot of programming, of course, for them. I guess most of them had mathematics degrees, but they weren't specifically doing anything necessarily in the so-called mathematical realm. They weren't doing numerical analysis, certainly, and in fact, I don't recall anybody that I met really doing numerical analysis on computers at that time. Certainly not in these data centers. Well, we'll come to that in a minute. I guess I came to numerical analysis totally by accident.

HAIGH: So in general, the systems analysts would design things, and they would give them to a separate group of programs to actually implement?

DUFF: Sometimes, yes. They would work together, actually, and there would be feedback. The systems analysts would run the show and the programmers did programming and got feedback. I was in that sort of loop, so I guess I was sort of a programmer. Because I was young, I was taken under their wing as almost like a kind of trainee. And at that time, this fellowship, the whole intention was that I would go on to become a mega manager at IBM, and of course I would have been stinking rich and living on 25 acres now if I had done that probably, but it was my politics that made me skip away from IBM.

HAIGH: All right, I'll ask you about it in a second then. So just one last question on how this was organized. The programmers, once they had written a program, would they punch it on some cards and drop it somewhere and wait to get the results back? Or would they actually be able to sign up for time on the computer to run it for themselves?

DUFF: No, no. In fact, the environment there was better than anything I had got until after I had graduated from Oxford. Basically there were card readers, and I think there seemed like a lot of people around to help then. I guess labor costs were less. So there would be operators around, but you could go to the card reader with your deck of cards, because that was always the main input in these data centers, and you could feed them in yourself. Then you would go to the printer and get your output and look at it. Or sometimes an operator would actually put the cards through for you and help you, but we were really hands on. And as I said, in the data center, at least the first experience, we could single-cycle machines, so there was obviously not a great pressure on the 1401. There's still not a whole lot of equipment, so they had non-electronic, more mechanical or electromechanical equivalent for doing some of the payroll runs. That's what my programming was sometimes about, moving them onto the electronic computer. Also, there's quite a bit of what they call nowadays "data mining". So they would want programs to analyze a file and say, "Tell me all the people who are over the age of 50 working for this organization." That kind of question was quite common.

HAIGH: That's what you would be using the RPG for, I assume?

DUFF: Yes, that's right. All of the data management in this language, RPG, yes.

HAIGH: So as I'm moving back to your own career, you had these three summers working for IBM.

DUFF: I reckon it was. I had one summer of looking for hut circles in the north of Scotland for some archeologist. It must have been the first one.

HAIGH: Then also you have been progressing with your studies, picking up this distinguished record as a mathematician. It says that you won the Mackay Smith Prize for the Most Distinguished Candidate in Natural Philosophy.

DUFF: Yes, I was a bit surprised with that, really, because I was thinking of myself more as a mathematician than as a physicist, but clearly I did very well on my physics final exams. I also got a prize for math, but that wasn't the first prize. I think that was the second or third, I don't know. I seem to always be proxima accessit, we call that. In school, in high school Philip Chalmers, the guy I mentioned, was the dux, and I was the second, the proxima accessit. So well, it's the same thing for maths in Glasgow I think. It may have even been third, I can't remember.

HAIGH: And that prize would be the Russell Bursary.

DUFF: Yes, that's right, that was the maths prize. It wasn't really a bursary in the normal sense. I didn't get continuing money from it. It was a lump sum, probably ten pounds. [Laughs]

HAIGH: So as you approached graduation with this distinguished academic record and this experience of working with computers, what kind of options were you considering for what to do next?

DUFF: So I decided maybe at the beginning of fourth year that I would be quite interested in a research degree, which brings back the politics. You see, although I'm not and never have been a Christian in the normal sense, I was involved with a group called the Student Christian Movement, or the SCM as it preferred to be called—it tended to downplay the "Christian" word at that time. This was a social one that was trying to get equal rights for black people, a lot of causes that are still not properly developed and won. So I was involved heavily with that. And part of it was against big corporations. So although I took IBM's money, and didn't ever think of not doing it, and although I worked for them in the summer and appreciated it, I more or less decided that it was immoral to work for a big corporation. So then a lot of other options were immoral as well. So it seemed to me that the only moral thing to do was to continue studying and do a Ph.D.. In addition, I quite wanted to do that. I like an intellectual challenge, and I thought it would be interesting to continue studies beyond the undergraduate level. So the combination of these led me fairly early to decide to do a Ph.D..

Then it was a question of the topic. I guess I always wanted to do maths rather than physics. Physics was an interesting subject, but there was a lot of lab work, and I didn't really enjoy the lab work so much. Maybe because that was always conducted in the afternoons after we had been to a pub for lunch and it wasn't always the best time of day. But one thing or another, I was always inclined towards mathematics. I then had to decide, of course, where to do the research. You don't just do maths, you have to choose a supervisor and a place to do it and everything else. I had a long talk with Robert Rankin, the head of the department, who said I was obviously a reasonable student because of getting prizes and all that. But I never thought of myself as being reasonable, and there were people that I think really were probably more natural mathematicians. Or to put it another way, that they couldn't do anything else other than maths, so they concentrated in maths and did very well. Those people, you would call them geeks nowadays, I guess. I certainly wasn't. I would like to think I've never been a geek, although my children sometimes think I am.

But anyway, it was a question of what to do. There was someone called Hunter, who I think is still involved in mathematics teaching in Scotland. Maybe he's just retired. He was very keen and did a course on number theory. In these days, "number theory" and "numerical analysis" sounded similar to me and in fact, since I've been editor of the journal, people send me number theory papers even though it's a numerical analysis journal, so there's obviously confusion somewhere there. But they couldn't be further apart, really. So I thought, "All right, let's do number theory." Number theory is fascinating, of course, all the very basic and elementary theorems in fact have become powerful nowadays with cryptography and that kind of thing. So I talked to Hunter and I talked to Rankin, and Rankin said to me, "I've got a friend down in Oxford who is one of the top number theorists in the country, he is called Brian Birch. I'll put you in touch with him." I can't exactly remember the sequence. I think I went down for an interview, although I'm not 100% sure. One way or another I was in touch with Brian Birch, and he said, "Sure, come along and be a D. Phil. student."

They had funding for me, but I actually applied for a Carnegie Fellowship, which was essentially restricted to Scottish people. It's from Andrew Carnegie, if you know of Carnegie-Mellon University in Pittsburgh, the philanthropist and extremely rich magnate in steel and such. So he had quite a lot of fellowships in Scotland. I'm not sure if mine was for Scottish people going to English universities or just Scottish people doing post-graduate degrees. But anyway, I got one of these, and that funded me, which was of course a delight to Oxford because they could then use their D. Phil. money for somebody else, but I was accepted without it, and I got that.

So I went down to Oxford in the autumn of 1969, originally to work with Brian Birch.

HAIGH: Now, I should say particularly for Americans who may be reading the transcript, my understanding is that in Britain, unlike the United States and I think many other countries, you would basically turn up one day and start working on your dissertation, and three years later you would hand it in. You wouldn't have this long period of taught courses that Americans institutionalized.

DUFF: It was nearly that case. Oxford has changed the system. Even during the time I was there they changed the system. At the time I went there, they did have what you might have thought of as a preliminary year, qualifying year, whatever, which was called the Diploma of Advanced Mathematics. You do this in three subjects. You do half time on the subject you are intending to do for the D. Phil. (D. Phil. is the Oxford name for Ph.D) and then a quarter time on the two other subjects. I was very much interested in computing, and not really realizing that number theory was something different I took that. Well there is computing in it, but it's not the same. So I did number theory, and then I did numerical analysis, because it sounded similar, as my half-time subject, and then I did the theory of computing languages, lambda calculus and that kind of thing, fairly theoretical. There was a fairly famous guy called [Christopher] Strachey, and he was my lecturer in that.

It was really an interesting combination to do these three, and it didn't take me very long in Oxford to discover that I didn't really want to do number theory. It was heavy in abstract mathematics. It was very heavy in algebra. Now, I mentioned to you that Glasgow was famous for analysis, as was Rankin, but Glasgow, apart from linear algebra, was not very strong in algebra at that time. So I came and I was meeting with undergraduates coming up from Oxford who had been through the system there, and I was absolutely swamped. It was the first time I felt out of my depth in maths, just with the heaviness of all the algebra there. I coped with it, I guess, and I was able to get myself up to speed. At the end of the year you do an exam, and I passed that quite easily. Martin Huxley was one of our lecturers. (Martin Huxley was a relative of the famous Aldous Huxley. I think he was a nephew). He was amazing because there was only one person that could follow the lectures in analytic number theory, and he was called Don Zagier. I'm not sure if he's got a Fields medal yet, but he's certainly one of the top mathematicians in the world today. He's Swiss, was younger than me, and he was doing his second Ph.D.. He was one of these semi-geniuses and still is. He was the only person who followed these lectures fully.

Anyway, that was the courses I was doing early on, and I decided around about the Christmas of that year, 1969, to check and see if I could switch, and do my main subject as numerical analysis and then have theory of programming languages and number theory as the two subsidiaries. Very instrumental in that, and this is probably important for our historical record, was the fact that Leslie Fox, who was then in charge of the Computing Lab, gave some of the lectures too.

HAIGH: So would you say it's an equivalent to the Master's degree?

DUFF: Yes. Well in fact, they renamed it a Master's degree. It's exactly the same thing as an M.Sc..

HAIGH: Oxford likes to have its own names for things.

DUFF: Yes, and in fact, we actually retrospectively tried to say we should be able to have an M.Sc., and the council blocked it. So it was still a diploma in advanced maths. It was almost the only year it was run. The next year people signed up for the diploma of advanced maths and were by mistake given M.Sc.s. So that's why we argued.

[Tape 1, Side B]

HAIGH: You mentioned a couple of well-known names there, Strachey and Fox, so it might just be interesting for the record if you could describe your impressions of them both.

DUFF: Well, Strachey, I don't remember as well. I hadn't realized Strachey was so famous a person until afterwards. In fact, probably until after he died, really. But he was a very energetic and entertaining lecturer. I actually did quite enjoy his lectures, there's no question. We had him, and we had a guy called Dummet in logic. He was a terrible lecturer, actually, but he was also very energetic. I remember them quite well, and I really thought they were good people and they did seem to be very bright people. I didn't appreciate their international standing at that time. Fox was really good. First of all, his politics were in line with mine. It was very clear; he wore his socialism on his sleeve. And he was really, let's say, a father figure. He was sort of—I mean this in a totally asexual way—a cuddly sort of guy, he was just a nice fellow. He was very enthusiastic in his subject, and he spread that enthusiasm to the students, including myself. In his first lecture—it was possibly a bit naughty of him—he really emphasized the fact that this was a topic for which there was so many applications, you would have no trouble getting a job, and a job you liked. You'd have no trouble continuing to do it in research. He said it was really a topic for the future. He was very clear on that.

I was living at that time in a room at the Maths Institute, which is where all the pure mathematicians were, including the number theorists. This was until December. The depression that was happening there was really quite sobering. There were people whom I knew to be good mathematicians who had just finished their Ph.D.s (or D. Phils) and were absolutely unable to get jobs, even though they were top, top, top-ranked people. They would just hang around the coffee time, no job at all, just living it up, probably from rich parents or something, but just hanging around there. Really, they were doing good mathematics still, but they weren't getting paid and nobody was employing them for it. There was a kind of depression there, there's no question. Whereas, of course, the upbeat nature of Fox and them made me think, "Well, as I've always said, there's more to my life than mathematics." There always has been, even when I was quite young. When I was at school, it was more then. So I very much decided and was won over by Fox. That was the sort of influence he had. He had a very broad knowledge of numerical analysis. He knew things: differential equations, linear algebra of course. He always had a very good eye for detail. If you wrote a paper, he could spot mistakes just by sort of glancing over it. He was very impressive at looking over stuff, and a really nice guy. I was quite won over to it, and I had no real hesitation about moving, and moved my office then to what was then the Numerical Analysis Computing Laboratory on Parks Road in Oxford. It's now part of Engineering; it's moved since. I never regretted that move.

Also, the people I met who were already doing the D. Phil., who were going to be my fellow students for three years, fellow Ph.D. students, were altogether in a different category from the pure mathematicians. They were really men of the world, in a sense, and were enthusiastic in the mathematics they were doing, but had a wider perspective, I would say. That was really something that I liked. I didn't meet the students, really, until I made the decision to switch, but that reinforced it. I certainly never regretted doing that change, which was a more radical one than I had realized. When I started, I thought, "Number theory and numerical analysis!"

HAIGH: They've both got numbers in.

DUFF: Yup!

HAIGH: So your Diploma in Advanced Mathematics thesis, or whatever it was be called, was "Network Analysis and Graph Theory". That sounds more like computer science theory or OR than numerical analysis.

DUFF: Well, I was influenced, you see, a little bit by the theory of programming languages, although I didn't put anything in there. I hadn't really settled down into my topic then, so there was actually at the end of the day quite a lot of sparse matrices in this because graph theory involves that. But the network analysis was very interesting at the time. People in Warwick were working on that and I went up there and visited and I got some papers from them. This was a very big topic in more operations research, I would say.

HAIGH: Is that network analysis in the sense that would have grown out of trying to analyze electrical power networks?

DUFF: There was some of that, and some flow networks as well, any network that has got a flow in it, whether the flow is of electrons or water. Graph theory is a very natural way of looking at that. I remember meeting Frank Harary at a meeting in Colorado and getting on quite well with him, and he was one of the main people in the area of graph theory. And Robin Wilson, of course, at Oxford then; I think he's at the Open University now. He gave lectures in graph theory in the diploma course. That was very good. That made me quite enthusiastic as well.

HAIGH: So did your work on the diploma involve programming?

DUFF: That's a good question. I can answer that tomorrow because I can look at my diploma thesis, which is there. I think very little. Wait! No, I must have done some…. The first term, until December, I never saw a computer. I think I did see computers when I moved to the Numerical Analysis department because, of course, it's rather important there, and certainly from the moment I started in the D. Phil. course I was quite heavily involved in computing, initially with a very famous machine called the KDF-9, which was located in our building in Oxford. The only input and output from it was paper tape. There's great stories about that, coming in one weekend and not being able to open the door. The reason you couldn't open the door was there was a guy called John Rollet who was a crystallographer and a very heavy person in computing. Alas, he died a little early because of his heart, when he must have been in his 50s, maybe about ten years ago now. Anyway, right up to the moment he died he was very active. He was generating great volumes of output on this paper tape. He had a big run going over the weekend. It had spewed out so much paper tape it filled up all the corridors, and was actually pushing against the door and it was hard to get the door open, and of course, you didn't really want to trample his paper tape.

I don't know what you did with that amount of output, to be honest, but once you've got your output on paper tape, if you run it through Flexowriter machines and print off stuff from it. Or you could read the holes in the tape, which some people were very good at doing. In fact, some people were not only really good at reading holes in the tape, but they could patch the tape, put in new holes for feeding and for input.

HAIGH: So that would be the standard five-channel paper tape, would it?

DUFF: Seven channel, actually. But it was reasonably standard for KDF-9s. There was no card reader for the machine, and there was no direct printer on the machine.

HAIGH: This was the numerical analysis computer, then?

DUFF: Yes, it was very much just us really, I think. As far as I know, nobody else was really using it at that time.

HAIGH: So did you have any contact with the main Oxford computer lab?

DUFF: Yes. What happened was, during the time I was doing my D. Phil (and I can't remember exactly when), they got some ICL equipment over in the main computing lab, which was across the other side of Banbury Road. Then we're coming to the Brian Ford days and Linda Hayes and other people who were involved in computing. Anyway, there's lots of people whom Brian, I'm sure, mentioned in his interview. And that was roughly the time they started to set up the NAG Library, which was called the Nottingham Group initially, because Brian was from Nottingham as a computational chemist. So he started in the computing lab, which was across from Banbury Road about that time, this is early 1970, '71 or so, and I used the machines over there. The one I remember the most was a 1906A machine from ICL. That had paper input, punch card input. We used to send our cards over, and they were fed through. The card readers had a habit of shuffling cards and doing other nasty things. It was quite fraught, really.

Then, at one time, this is rather famous, they had decided that the punch cards were too expensive in the budget. We were using a language called ALGOL 60 (we used ALGOL 60 and ALGOL 68, but I think it was ALGOL 60) at the time. You can separate the statements with semicolons, so you could have several statements on one card, unlike FORTRAN. So they would give you as a postgrad student a certain set number of cards, and you had to beg, borrow, or steal from your colleagues extra cards if you needed them. So what you did then was to compress these and have three or four statements on a card, which made it totally unreadable, of course. This was an amusing thing, I think. People have restrictions on resources, and still do, but this was a card restriction we had then.

Oxford computing had a tradition of Ferranti machines. It became part of the British company called ICL, which was taken over by Fujitsu I guess and then eventually sort of went to the wall.

HAIGH: I think they have some pieces left; they just rebranded it all as Fujitsu. They decided the ICL brand wasn't worth keeping alive.

DUFF: Right, yes. They were quite big at that time, and of course, they were pushing themselves as being a British company, so the universities were being obliged to buy their equipment, even if it was inferior to IBM's at the time. Which was one reason why I started at Harwell, because the computing at Oxfordwas really quite bad, to be honest. The people working there were fine, and I'm not complaining about any individuals, but the system wasn't great, and this card rationing at Oxford is one good example.

It's very interesting to know how I first became involved. I think Leslie Fox felt that the research group at Harwell worked in my area of expertise. But it was mainly Alan Curtis and John Reid who were at Harwell at the time doing sparse matrices, and I got to know about their work, of course, and I went down to visit them. They were quite excited about the prospect—I don't think they had one before—of having an Oxford D. Phil. student working with them. So I then used to go down from Oxford about one day a week on my little Yamaha 50cc cycle or by bus to Harwell, which seemed like an awful long journey these days. I think I visited more often towards the end of my D. Phil.. But the main thing was that I did most of my computing there, and they had all IBM equipment. It was chalk and cheese. That just worked fantastically. You could do a lot of computing. The issue was, I think they had an ALGOL compiler, but really they were a FORTRAN lab, so I always had some FORTRAN, I had even done a little bit back in the data center days. So I wrote in FORTRAN, I wouldn't say most of my computing; it was less than half, maybe. But about a third was done on the machines there. Certainly, some of the heavier stuff was done down at Harwell.

HAIGH: So that relationship with Harwell is something that you initiated? It wasn't a regular practice?

DUFF: No, at that time it wasn't. There's been some students since then. There was a guy, Pat Gaffney, who was working with Mike Powell as a supervisor, and I think he was about contemporaneous with me in Oxford. He may have been a year after or the same year. He was also being partially supervised down at Harwell. Alan Curtis in particular, was the vocal guy that you noticed there. He just celebrated his 80^{th} birthday a couple of years ago and is still doing okay, and still intellectually active. He should not be confused with the NPL Alan Curtis, I should say. So Alan Curtis was in charge of the maths group at that time, and in the members of the maths group were Mike Powell (who later went to be a professor at Cambridge and became an FRS) and John Reid (who works for me now as a consultant). Roger Fletcher, who had been working with Powell on many different algorithms (Fletcher and Powell algorithms) had just left to go to Dundee, where he has just retired recently. I wasn't there to replace him because I don't work in the same area, and Roger was more senior to me, but eventually when I got the job there, it was sort of his job. But Roger was there when I was still a student.

The other great thing is, they had lots of visitors. People like John Dennis from America and his student Bobby Schnabel were over there, Don Goldfarb from New York, they were all coming to visit. Mainly to visit Mike Powell, I guess. Alan didn't tend to have a lot of visitors other than John at that time, but maybe John got a couple. It was a very active place for numerical analysis. It was the most active place in the UK at the time, without any doubt. So it was really an exciting place to be. Alan Curtis came to the seminars in Oxford, which we were more or less obliged to go along to as graduate students, even if we didn't understand much at the time. And he has become a little bit calmer as he got older; certainly now, but even 20 years ago he was a bit calmer than he was 30 years ago. He used to tear them apart, some of the people, especially if they were overstating things or trying to pretend that they knew something that they didn't. This was great sport for us undergraduates. It was like being at the Roman circuses. I really loved him and his influence. So I suspect it was a combination of Fox, Curtis, Reid, and Powell that somehow got me down there.

What happened was that I did the diploma in advanced math, and I was very much under Fox' tutelage at that time. I didn't really have any much to do with the other people in numerical analysis or with Harwell at that time at all.

HAIGH: And was Fox…

DUFF: He was my supervisor then, if you like.

HAIGH: And was he the most senior person in the numerical analysis?

DUFF: Oh yes, he was. Very much so. He was the only professor. These were the days when you only had a single professor in the department, and he was the head of the computing lab. He was Strachey's boss as well at that time. He and Wilkinson were probably the two most senior figures in British numerical analysis at that time. Of course, Wilkinson never really got involved in teaching. He was researching at NPL, all the time. He got involved in teaching at Stanford later, but not in the UK. So I was very much involved with Fox in that first year.

In the second year there was a visitor that came from the State University of New York at Stony Brook in Long Island in the USA. He was called Reginald Tewarson, and he was a character, fairly flamboyant, nice fellow. I am still in contact with him. He was just finishing or writing a book called *Sparse Matrices*, so what better than to team up with him? So he was effectively my supervisor for the second year at Oxford, which was the first year of my D. Phil.. But I finished the D. Phil and the diploma in three years, so it was the middle year in Oxford. But he was only there for a year on sabbatical, and when he left Fox had to do something with me. Fox decided that my knowledge in this area was… well, he felt that he wasn't the appropriate supervisor for that particular specialty. So I think it was at that point that he decided, "Well, there's a group at Harwell, and what they're doing is sparse matrices!" It was Alan Curtis that got it going because of his need for sparse matrices for solving stiff differential equations using Gear's backward difference formula. So that was the origins of the sparse matrix work at Harwell, and I would say that John Reid was the main person—well, he and Alan Curtis were the two people working on sparse matrices at that time. They had a few papers from about 1972 connected with that. So it was very natural for me to go there. I worked with both of them. Fox was still my nominal supervisor all the way through and I worked quite closely with John Reid in my third year in the D. Phil.. I went to Harwell at least once a week.

HAIGH: All right. Let's get some context on Harwell at this point then. You've mentioned the strengths of the numerical analysis team there. How about the institution as a whole: What was the place of the computer people within the overall framework? What was the general mood and culture of the place?

DUFF: Well, it was very exciting because, of course, there weren't any personal computers or laptops in those days, so computing was by its very nature more centralized. It had to be. Alan Curtis in particular was heavily involved even in developing operating systems for the computers, as well as mathematics. This is what people in mathematics were doing these days—they were jacks of all trades, and mastered them nearly all as well!

So he was involved with a thing called HUW, which was the Harwell User's Workshop, and this was one of the first non-card-entry systems in the world. We're talking here about the early 1970s. So I did use card input, but I could also use teletype input, more or less sending stuff directly to the computer.

HAIGH: So that was a time-sharing operating system, was it, or just a remote batch-entry?

DUFF: I used it as batch-entry. I don't think it was time-sharing. But it was really very innovative. For its time, it worked extremely well. In fact, I went to US government labs years later in 1976 to NCAR for example, and it was primitive compared with that, still using card input to a Cray 1 machine. So HUW was really quite advanced.

Harwell was experimental, very theoretical as well, and computational is now recognized as a separate field. I would say that this is one of the first places in the world to have all three branches of science working together. The computational people were theoretical physicists who wanted to compute things, experimental people who wanted to analyze data. They came mostly to Alan, who's a very underrated person, I should say, an incredibly talented, incredibly good mathematician. So he coordinated these things. Around about the time I eventually came to Harwell, the numerical analysis group was set up under Mike Powell, and maybe just a bit before then, which had Fletcher and Reid and Mike Hebden as members of it. Sid Marlow and Mike Hopper were also there in the group.

HAIGH: So the numerical analysis people were very closely linked to the main computer service center?

DUFF: Absolutely. We were actually part, at that time, of the theoretical physics division. It was about the time I arrived at Harwell, maybe the year before, but certainly after I came to Harwell as a graduate student. When I was a graduate student, I was in the theoretical physics division, and that was where the computer was. That was where all the numerical analysts were and this main group of applied mathematicians, maybe not all of them. It was people like Jack Howlett. Bill Morton was involved in it. It was just before my time that they were there at Harwell. They were the people who started using and supporting users on the computers for the first time at Harwell. Howlett, and Bill was involved, and these other people that I mentioned. Bill Morton became a professor at Oxford eventually. He was a junior professor at that time, I guess. Howlett was the senior person.

HAIGH: Harwell's main mission was nuclear power research, wasn't it?

DUFF: Yes, it was called AERE Harwell, Atomic Energy Research Establishment. But of course, that involved an awful lot of different sorts of physics and chemistry because materials were a big thing, a very big thing, metallurgy, because with nuclear reactors it's very important that they don't disintegrate. So there was a lot of materials physics, and finite element analysis eventually became a big thing when computing took over. So there were a lot of different kinds of physics and chemistry involved, although the main driving reason for it was the nuclear industry. By the time I got there, it was purely nuclear power. When it was set up originally, it was nuclear power and weapons, and in fact Klaus Fuchs, the famous atom spy, took his secrets in a briefcase through the Harwell main gate and gave them to his handler. There isn't a plaque to that, to be honest, but he was working there at the time.

The Atomic Weapons Research Establishment, AWRE, was set up some years before I joined Harwell. I can't remember exactly when. Just before I got there, the computing for Harwell was actually handled by AWRE, and there the turnaround was twice a day, because you put your cards on a bus or a van and went down to AWRE, ran it, got the output, and brought it back. It came back during the day, so you had two runs a day. That was not long before I came. But by the time I came we had our own machines, and we were pretty independent from AWRE, and certainly there was no nuclear weapons research at Harwell at any time, which would have caused me a problem because of my politics again. Nuclear power I've never had a problem with, and I still believe it's the main option for power in the world.

HAIGH: That's interesting. So you had been turned away from IBM by the moral objection to a large corporation, but at this point…

DUFF: But Harwell was like a super university. It was fantastic. They had really some of the best brains in theoretical physics, metallurgy and applied chemistry interalia probably in the world, but certainly in the country.

HAIGH: So was it only in the 1970s that leftists began turning against nuclear power?

DUFF: Yes, well, leftists, yes, well. I think there are some misguided people, "Friends of the Earth" sorts of people. They're the kind of beard-and-sandal brigade that are looking for issues that don't exist. They're not social reformers. They're not interested in the poor, not interested in the fact that Glasgow had more slums than Naples when I was a boy. That's the kind of thing that's important. So yes, I don't regard them as leftists. But yes, it was depressing to me, the fact that people could be against nuclear power the way they were at any stage. Now, I'm not saying that there aren't issues about nuclear waste, and they're issues that to some extent have been addressed and to some extent haven't, and still to this day it's the case. But when you compare it with other forms of energy generation, you see now how important it is, having it as one of the options. I'm not saying it should be the only option, of course. So I was never in any doubt that nuclear power was good, so I don't have any moral qualms about working there. I would have done if they had weapons as well.

HAIGH: A couple of follow-up questions there on your first impressions of Harwell, and then we'll return and wrap up your Ph.D.. So you mentioned that the computer center has IBM equipment installed on site by the time you came in and it worked much better. Can you remember what kind of machines that they had?

DUFF: In the'70s by the time I actually came to Harwell towards the end of my D. Phil.. Well, I didn't go straight to Harwell after my D. Phil. anyway. It was IBM 360s and 370s. It was big IBM mainframes. It was what you call "big iron".

HAIGH: So they were pretty large machines?

DUFF: Oh yes. And they were the kind of machines that you would find in data processing centers. They were the front range big IBM machines. They did have an IBM Stretch at one stage back before my time, and that was really more designed for numerical computing as opposed to data processing. They had the TSO operating system, and MVS. They were just diabolical to use, really, from my point of view. I mean, I fell out with IBM (not with IBM personally but IBM machines) later when I came there because of the operating systems. They're just so big and cumbersome; they were not designed for numerical computation. But we were using machines that were general machines. Big iron. Again, I've got plenty of notes on these machines somewhere, but my memory is not good. Some people are much better at that, but I am not.

HAIGH: As you later had an opportunity to see the American DOE Labs. How would you compare Harwell with, for example Argonne?

DUFF: Well, of course, I would eventually spend a whole year at Argonne, some years afterwards, but I would say that if you're going to draw parallels in the US, then I think Argonne would be the closest. I've been to all the labs; I've been to the semi-closed one like Los Alamos and Livermore, Berkeley, and Oak Ridge. Quite a lot to Oak Ridge, actually; early on, in fact, I was offered a job there. So I would say that Argonne was the closest in my mind to what Harwell was like. Again, one of its main reasons at one stage was nuclear power in conjunction with Fermilab. It diversified, of course, as Harwell did. But it was quite a big core mission. And the maths and computer science division at Argonne was pretty central and at one stage did a lot of computing for other people in the days before people had their own workstations. So there were similarities. But of course, when I came to Argonne I was more senior, so I wasn't as impressionable. Certainly when I was at Harwell, I was very much impressed by the quality of people. I'm not saying that they were necessarily better than Argonne, but I was more aware of how good people were there. That's what makes my heart bleed for what's happened to the place now.

HAIGH: At this point, in the early '70s, do you think that the level of support Harwell was receiving and the kind of facilities that it had would be comparable with the American labs?

DUFF: Right, I would think so. Well, as I said, my experience when I first went over and did some computing in the mid-'70s in the US was that we were maybe even ahead of them slightly. Certainly, I would say it was commensurate with those labs at that time, but I wouldn't know the detail. When I was a student, I didn't know what the US labs could do.

HAIGH: Let's leave Harwell on one side for a moment now and return your Ph.D. topic. You mentioned that it was related to sparse systems and to the work that was being done at Harwell. So could you talk a bit more about what you were studying and what you found out?

DUFF: Well, it came about quite naturally after my dissertation on network analysis and graph theory, so I was very keen on that side of the sparse matrix world, which is really to say what's called direct methods for sparse matrices. It was really methods based on a rather simple algorithm of Gaussian elimination, and I've really been working quite a lot in that area since—quite amazing after all these years.

One of the most influential things, I guess in my whole life to be honest, was a visit I made to IBM in Yorktown Heights around 1971. Ralph Willoughby was there (again, people die off, unfortunately, he would have been a great guy for you to interview) who was in the mathematics section along with some really very energetic people: Gary Hachtel, Bob Brayton, and Fred Gustavson, who was a young programmer at that time. That group were really energetic and working very much at the forefront of sparse matrices. I went there because there was an international meeting on sparse matrices. In 1971 it was the second one to have been organized in the whole world. The previous had also been at Yorktown Heights in 1968, and I had seen the proceedings from that, but it was before my time. I went there and I helped Rose and Willoughby, who were doing the proceedings for that, by doing proofreading for the papers as they came in, because I stayed on at IBM for the summer or part of the summer after that meeting. [Rose and Willoughby (eds.): Symposium on Sparse Matrices and Their Applications. New York: Plenum Press 1972, pp. 157-166]. That was extremely influential. That was really what set me on my course I would say, more than anything, more than Harwell even, really. It was the enthusiasm that these people had. Ralph I've always looked at as being the father figure, or perhaps a grandfather figure, in sparse matrix work. He made lots of pronouncements. One was that he thought there was a few more years left in sparse matrices, it was good for the D. Phil., and I could get maybe ten years research out of it. I think he was a little wrong there. I think it's still a very energetic field.

HAIGH: You thought that at that point that all the problems would have been solved and people would have to find something new to think about?

DUFF: Yes, I guess so. That's right. In fact, we were solving problems so fast you couldn't count at that point. A lot of them were connected with graph theory, a lot of basic stuff. Donald Rose was one of the people and he had edited with Willoughby the proceedings I did some proofreading for. Of course, by doing that, I learned quite a lot. So I came back really energized from that, and that really drove me through, along with the support I had at Harwell and very good support from people like John Reid and Alan Curtis. So Yorktown Heights is still a very active research lab, but not so important for sparse matrices. Fred Gustavson is still there, but not too many other people.

HAIGH: So it's really that meeting that introduced you to the community of people working on sparse matrices.

DUFF: Well, yes. I met a young fellow called Alan George for the first time, who later became professor at Waterloo and is one of the more significant figures in our field. I met several people there with whom I've had contact ever since. That was very important for becoming known in the States. I've been going back to America ever since, initially once or twice a year, and now four or five times a year. For one reason or another, there's more people in the field there and it's very much stronger in the US than it is here.

HAIGH: So how well developed was this sparse matrix field in 1972? Was it a technique that, in way or another, people had been using since the early days of computing, or was it something that had just been getting documented a few years prior to the meeting?

DUFF: Well, various so-called relaxation methods were developed by Fox and his predecessors, and these led to modern iterative schemes with people like David Young and later on Dick Varga. These were quite well-established by the time I was a student, including successive over-relaxation, and various forms of iterative techniques were very popular. And these were the ways that people solved large problems, mainly from discretizations of partial differential equations. But sparse direct methods were absolutely not at all established until pioneering work that I had mentioned at Harwell, with Curtis and Reid, and at places like Yorktown Heights. And one of the main driving areas was not PDEs, it was actually ordinary differential equations with Gear's backward difference method that I mentioned where the sparse Jacobians could well be asymmetric and very sparse. That was really where the initial impetus came for direct sparse matrices because iterative methods did not have a chance there, and still don't have much of a chance in that area, so it had to be direct methods. And of course, then a large sparse matrix was 100x100 or 200x200, and now it's a million by a million in the context of direct methods. So it's really amazing how things have changed.

HAIGH: Would the matrix itself in these cases be stored as a regular array with a bunch of empty spots in it, or is it stored in a data structure that would be more efficient.

DUFF: That was exactly what was happening in the research field at the time of 1970. To store the whole matrix, even 100x100 in these days, was too much for the machines they had. They didn't have enough memory. So they used things like bitmaps where they had entries stored as single bits to identify where the non-zeros were, and then they stored the values separately. But that was rather slow and cumbersome. So then they developed standard techniques for just storing the non-zero entries of the matrix, which are in use today. They haven't really changed from what was being developed at that time.

HAIGH: But in the very early days, the problem was not so much the storage as just finding a method that would work in some vaguely acceptable time period with a relatively large matrix?

DUFF: Yes, it was really quite a big programming effort. What we were trying to do really initially was to look at methods for dense matrices and then translate them into this sparse environment with a sparse data structure. That was really some of the origins of direct sparse matrix work, and since then it has developed rather a lot.

HAIGH: I have more questions on sparse matrices, but I can return to them later if you think we've dealt with them enough now to cover your dissertation work.

DUFF: Yes, I guess so. As I said, I'm glad that we brought up the IBM thing because that was quite crucial and just came at the right time. I must say that I think it's a wonderful thing for graduate students to be able to travel. I think the lab actually funded me. It cost me 80 pounds to fly over at the time, and they funded that. That was Fox from the computing lab. I think those sort of chances are so important at that time in someone's career.

HAIGH: And that was your first trip to the United States, was it?

DUFF: No, it wasn't. It was my first trip working-wise. I went to Toronto Data Center with IBM, if you remember, in 1967. That was my first trip. Then I visited the United States when I was there. I came through Niagara and visited other places.

HAIGH: So you've mentioned that conference. Do you want to say anything about the actual content of your D. Phil. thesis?

DUFF: Yes, I'll show you it tomorrow, the D. Phil. It's probably got the record of being the thickest D. Phil to come out of Oxford, but then I cheated a bit and had some program listings in the back. But it was a general one. I would have said that it has more breadth than depth as, although there must have been enough depth for it to be a proper D. Phil.. But I was really trying to cover a wide area. In fact, on the strengths of that, I was invited, I think that Ralph Willoughby made the recommendation, to write a major survey article for the proceedings of IEEE, which is still my most quoted and famous paper, which I finished in 1977. Really, it was on the strength of having a very broad view of things in my D. Phil. that that came about. There were three or four published papers that came directly out of the D. Phil., maybe slightly more eventually, so there was certainly enough depth as well, but I was fairly generalist in covering things. I was trying to hedge my bets at the time, because it was a fairly new field, so who was to know which was going to be the real rich veins to follow? I certainly didn't know at that time, so I kept a broad brush.

Also, I was still very involved with lots of societies. I was president of our Middle Common Room at the last year at Oxford.

HAIGH: And what does a Middle Common Room do?

DUFF: That looks after the welfare of graduate students, mainly.

HAIGH: So it's "middle" in that it's halfway between senior and junior common rooms.

DUFF: That's about right. It's probably a lot closer to junior than senior, because they've got all the power. But we were represented on the governing boards of the college, and you have to make representations for the graduate students. I had quite a lot to do then because they were changing the charging structures for the college, and what I made them realize is that graduate students are essentially 52 weeks a year people, and not just the same as undergraduates with three eight-week terms. So we managed to negotiate quite different and better deals for the graduate students, which I think is a legacy that's there today in New College. I'm not sure it's in all the other colleges. But that was bitterly contested, and all my debating powers from my days at school had to come to the fore when I was addressing these very austere and seemingly very aged people from the college.

HAIGH: Now, while we're talking more generally, did you find that the general culture and atmosphere was very different at Oxford from what you had been used to at Glasgow?

DUFF: Totally. First of all, it was a sherry culture there in these days, and it took me not too long to discover that sherry is stronger than beer, but I think I behaved myself most of the time. Well, of course, for me it was a big difference because it was the first time I wasn't living at home. For all the time I was there, I lived in the graduate building of New College, called the Sacker Building in Longwall Street. That was very much our domain and my domain when I became president and we organized all the dances and parties and stuff. A lot of my main activities were outside of the university, because in these days there were only five ladies' colleges, and these girls were hotly contested, to try and go out with. Many of them were not my type of girl anyway, given my socialism and that.

So I got involved in a lot of completely extramural activities, and one of the main ones was called the Oxford United Nations International Group, which ran from the British Council in Beaumont Place in Oxford, and I became secretary of that. We used to run this on a Friday night. We did lots of things at weekends with students coming to learn English. Indeed, I met my wife through that who came along for the same reason, to help in the club—she was English. But I was very much involved in a lot of other organizations, like the SCM to some extent but that was a little too religious for my liking in Oxford; and JACARI, the Joint Action Committee Against Racial Intolerance. We used to do a lot of tutoring of people that had come from the Indian subcontinent to live in Oxford, with children who didn't speak much English. I used to do maths tutoring with them as well. So I was doing a lot of stuff. Those were university organizations, but I did quite a lot outside of that. And that's another reason why I was flitting about, trying to figure out what to do, so I was fairly general rather than getting my teeth totally down into anything.

HAIGH: All right. Well, maybe we should break there then. When we come back, I'll ask you about how you made your career plans on graduation, and what's happened next.

[Tape 2, Side A]

Beginning of Session Two, on the afternoon of the 31st of August, 2005.

HAIGH: To pick up where we left off, we've discussed your experiences at Oxford and your thesis. One topic I should raise at this point would be Jim Wilkinson, another very well known figure. So had you had much contact with him as you were working on your thesis?

DUFF: No, it probably came later. I got to know Jim probably at meetings like the Gatlinburg meeting, now called the Householder meetings. I got to know him reasonably well eventually. I think when I was a student I knew of him, of course, and I think I did meet him, but I would say I didn't really have much contact with him until afterwards.

HAIGH: All right. I'll ask you about him later, then. So as you approached graduation, what kind of options were you considering?

DUFF: I was reluctant to move into a business for the same reasons as earlier, so the natural option was to move into universities as a lecturer. So I was looking around although there were not too many jobs then. I think I only applied for the one that I got. Also, as a long shot I applied for a Harkness Fellowship, because as I mentioned I already had a very good experience in the US and knew that my field was big over there. Reggie Tewarson, of course, had been my supervisor and he was back in Long Island. So I applied for this Harkness Fellowship, which is, I must say (if it doesn't sound too immodest) an extremely prestigious fellowship, and probably the most prestigious in Britain at the time for spending two years in the US. I had three sets of interviews for that, including aptitude tests and debates. Of course, my old training helped me there. Most of the other candidates were public school boys, so I was able to out speak them quite easily, and I got this fellowship.

At the same time, I had been interviewed by Ewan Page, who was a significant figure in British computing, and in the hierarchy of universities because he later became vice chancellor at Reading. He interviewed me at Newcastle upon Tyne in Northeastern England, and offered me a job there at the same time I got this fellowship. So I had a long, agonizing debate and decided that the fellowship was too good. I could risk finding a job again when I came back, but I had to do that. So I told Newcastle I was turning down the job because of that, but they knew about this prestigious fellowship, they said, "Oh no, this is great! You do that, and then come join us. How would you fancy that?" Now, whether it was a good or bad decision, I don't know. I mean, I enjoyed my time in Newcastle. In fact, that meant that I had a year and a half in America and came back to a ready-made job in Britain. So both were satisfied at the same time.

HAIGH: So would this Harkness post-doctoral fellowship be tenable at anywhere in the United States?

DUFF: Yes. It gave a lot of money. It also obliged you to travel. In the 12 months you had to travel for at least two months to different regions of the United States from those in which you were working. It was partly to try and foster international relationships between Britain and the US. So one had to do that, and that wasn't a problem; I like traveling and doing things other than maths as well.

So I organized it, and having the contact of Tewarson probably helped me get it, even. So I started off going over to Stony Brook to work with Reggie Tewarson and other people there and met people like Stan Osher and other quite famous people who were there then. Also it was not too far from Yorktown Heights, so I also visited people there during the time I was at Stony Brook. After that, I did my tour. It was a tour on a US visit program, coordinated by the YMCA. So I stayed with lawyers, clergymen, doctors, those kind of people, traveling throughout the southern states. In 1972, and that was an interesting time to be traveling in the southern United States, because basically the slave trade was finished, I guess [laughs], but the slavery in terms of everybody having black servants hadn't changed. So you were very conscious of that. Also, there was a general distaste for New Yorkers, and my car had a New York number plate, so in western Virginia I had to be rather careful and speak with my broadest Scottish accent immediately before I got gunned down with a shotgun. So it was quite, I wouldn't say wild west, it was wild south. And it was a really interesting time to travel. The Vietnam War was not quite finished at that time. Nixon had just taken over and either negotiated or was kicked out Vietnam the year after I was there. So I actually went to Washington on a demo with about a million other people, which was quite exciting, and the culmination of my student experience for demonstrations.

So I did this tour, which was great, throughout the southern states, met my then fiancée, future wife, a couple of times during my year there, and carried on the second half of the Harkness in Stanford. That was where I first got to know Gene Golub, who has been a very big influence in my career since then. And of course, the Stanford body of students included many people, like Margaret Wright, who has now become quite famous in her own right, and Walter Murray who had come over from England before. So there was a great group of people there, both students and faculty. I spent a term at Stanford, and came back for this job in Newcastle. So in a sense it was bad because I was really getting on well and I could have stayed on longer. The fellowship would have let me stay another nine months. I could have stayed longer, and ideally it might have been good to have done so, but I had this concrete job and they weren't really very keen to hold it over any longer, so I think it was in November that I came back to Newcastle.

HAIGH: So you wanted to talk a bit about Gene Golub at this point?

DUFF: Well, yes, I'll be talking about him quite a lot later, but that was the first time I met him and the thing that I was impressed about then, and very strongly since, is just how very well-disposed he is to young researchers and how much he tries to help them and encourage them and give good mathematical advice, ideas for doing research, that type of thing, just exceptionally good. I mean, he's got an ability to talk to young people that not many senior professors have I would say, even to this day. Certainly in these days, since I think professors are better now than they used to be. He was very encouraging to me and my career when I was there back in 1973. Of course, he was traveling as he always does, and so I didn't see him a lot of him when I was there. My interaction was more through his secretaries and students and colleagues.

But there was great activity there. Serra House which was the building where numerical analysis was, was an old building on campus that's long since demolished. Tony Chan was there, and some people who have become pretty famous in America, just a little younger than me because they were still doing their Ph.D., and I was a post-doc. But as you know in America, people take a long time to get their Ph.D.. Margaret Wright, for example, is older than me, and Mike Saunders and Alan George had only just left Stanford at that time. I had met him already at Yorktown, of course.

The Stony Brook visit was excellent. When I was there I managed to finish off some papers from my thesis, which is what one always wants to do first of all as a post-doc. I had some interaction with Reggie and other people at Yorktown Heights. So that was a good period, but the Stanford visit was really quite exciting.

HAIGH: Did you do any teaching at Stony Brook?

DUFF: No, basically not. I'm not saying the fellowship forbade it, but it was quite well funded. The fellowship actually gave money to the university as opposed to the opposite. I gave some seminars on my research work as I did at Stanford, but no teaching at all, not to undergrads.

HAIGH: So you only mentioned the circumstances by which you came to take the job at Newcastle. What was the general state of the computing laboratory there at that point? Was that the same thing that became the computer science department, or was that separate?

DUFF: No, I think that would be separate. I'm not sure if there was a computer science department even then. I'm not sure. I thought it was still called the computing laboratory.

HAIGH: It's possible. I knew they have something that functionally is a computer science department. Of course, in some universities, the computer science department became a separate thing from the computer center.

DUFF: Right. So at that time, of course, there wasn't the distinction as you say, and I'm not sure exactly what has happened since. But basically the department, as many of the departments were, was founded by numerical analysts, to some extent Ewan Page, who was the professor and head of the department who was in, I guess, numerical computing. And then Brian Randall was the other professor, and Brian Randall is not at all in numerical; he's in computer science. They were the two main professors in the department at the time. And then there were other people who were not too much different in age from myself, a little older. Ken Wright was the senior numerical analyst there, he must be over 65, and I think is emeritus there now. I've seen him quite recently. There was quite a lot of core activity in numerical analysis there. But by this time, I was still working with the people at Harwell, mainly Curtis and Reid, so my research was with them. I hadn't really picked up so much direct collaboration with Americans at that time, so it was mainly at Harwell, and I was continuing to do that kind of research.

HAIGH: How did you find the experience over all?

DUFF: Well, I loved Newcastle. It's not so different from Glasgow in many ways. The people are extremely friendly. It's a lovely atmosphere to be in. Really genuine honest-to-goodness people, no pretensions, no trouble with social mix, no trouble with politics—it's all pretty left. I did settle in quite well, and I got married in Newcastle, or in Heddon on the Wall just outside Newcastle, and we had our reception in the university house out on the banks of the Tyne. So the whole thing was just great. On the weekends, I used go off to the Lake District hiking, out to the Northumberland hills, or into the south of Scotland or even travel further north. It was just a great environment, and I'm still good friends with the people in the department,—I was best man at the wedding of one of them; he was a chief usher at mine. We still have a lot of contact. They're in New Zealand now. It was quite a good mix. In those days, it was computer science and numerical analysis all in one department, and the computer scientists didn't know much numerical analysis; we didn't know much about operating systems, not the research side of them. But we were all great friends and might meet socially and talk quasi-work-wise, so it was very good.

It comes to why I moved from there, which is what we'll come to next. I enjoyed the teaching; I didn't enjoy the marking of scripts, especially when you did sorting algorithms and things that were almost impossible to debug. So some of the script marking, was terrible; but the lecturing, I quite enjoyed. And we had quite good students at Newcastle then, so it was fun. But what happened was that Roger Fletcher left Harwell and moved to Dundee. That had created a vacancy. Mike Powell, who was head of numerical analysis, this was now within the Computer Science Division, contacted me to say that there was a good chance of a job, and certainly would match my salary or do better. It was a great environment. I was agonizingly torn, really agonizingly torn. Paige was none too happy, didn't want to lose somebody he just spent a year waiting for. And this was in 1974—I had just been at Newcastle a year. I told Powell that it wouldn't be appropriate for me to leave immediately, so we eventually came to a compromise, and after two years I went down to Harwell, with Mike encouraging me. Also, Harwell really was for a while a very exciting place to be. In my field at that time, large-scale computation, it was much better. It was much more academic at Newcastle.

HAIGH: Did Newcastle have a large computer?

DUFF: We did. I wish I could remember the exact model, but it was, again, an IBM. Elizabeth Barraclough was in charge of it. Brian Randall, I think was a fellow at IBM or something; I think he went to work for them eventually, but even at that time he was well-supported by IBM. He's a great mate of Edsger Dijkstra, who came to Newcastle a couple of times when I was there. So there was quite a lot of excitement, a big machine there. On the numerical side in addition to Ken Wright there was John Lloyd, and another numerical analyst who died early – Brian Shaw if I recollect. So there was a little corps of us, and they're good people, but not in my field particularly, so I wasn't getting great stimulation there.

HAIGH: So you had mentioned that you were torn. What was what?

DUFF: Well, life, academic life. I liked Newcastle. I loved Newcastle, and in fact, in some senses, still do. I mean, I still think that's in some ways a better place than down here. It is idyllic here in this village, and I do like it, and I wouldn't now move, but I was really torn to come down south. I mean, the south of England is not something that a Scot readily goes to, whereas Newcastle is almost in Scotland. So there's a lot of other things: personal life—you know, I spent all that time in the Lake District, Northumberland, the Moors…. Well, around here, the hiking is relatively rubbish compared with that. It's not bad; you can get to like it as we have, but it's not the same. So it's a lot of reasons, but very much the work or research, well that's what's done it. It was clear to me that it was far better at Harwell, and I've got no doubt that I made that decision in that way. If I stayed in Newcastle, God knows what would have happened, really. I really don't know.

HAIGH: So would you say that, at that point, people within your group at Harwell would have had the same kind of freedom to set their own research agenda that they would enjoy in the academic world?

DUFF: More. Well, as much. Well, maybe more: the same time and freedom to do research—this was the great thing—without the chores of having to do teaching (well, if you call teaching a chore), without having to do marking, without any of this stuff. So it was, in that sense, far better than a university. Then, it was all the freedom of a university. Powell and Reid knew what I could do, and was interested in doing; that's why they recruited me, and I did what I wanted to do. But on a daily basis, on a precise basis, no, there's no control at all. You produce papers, you do good work, then fine. And applicable work. And that's been my philosophy since; at CERFACS, which we'll come to later, and as group leader myself now at Rutherford and Harwell. Whatever the source of money is, to try and defuse that from the research to some extent, so people are free to do what they're good at doing. That's why you recruit them for doing it. So you could not have a better atmosphere. You wouldn't have it to the same extent in the Department of Energy lab in the States, as I discovered. The Department of Energy labs, you've got this "soft money" aspect, and I've got this now in England a bit. You've got to put in for a grant, a big grant or a small grant, or money that pays your salary essentially, to the DOE. You're fairly free; but there's certainly a lot (and this sounds bad to say) of accountability. I think you should be accountable, but there it was more oppressive, I would say. So that was really the reason for being there and wanting to be there. And you got a lot of feedback, because physicists would come and they'd have problems on neutron diffusion or other things, massive problems which they knew enough to pose numerically, but they had little idea how to solve. So the combination was just great, really doing stuff that benefited and helped people and helped society as well. So there for a nice socialist person, this is good.

HAIGH: This idea that the work you would do would be something that would help the users within Harwell. Was that something that there was some kind of pressure towards doing, or was it just in your case was something that came naturally, whereas it might not have done for everybody else?

DUFF: It sort of came naturally to the kind of work that me and my close colleagues were working in. The accountability, I don't know if the word was in the English language then, but it was quite different; it really was. You were judged for promotion or your pay or whatever on the kind of scientific output—papers published, talks given, professional and technical things of that kind. There was actually very little exact pressure on accountability in terms of delivery to people. But all the time, it was in our blood. We recognized that numerical analysis can really benefit other people. That's why some of my colleagues and friends who are in academic life in numerical analysis (who shall be totally nameless in this interview) are very well-regarded, but I think their work is not that great because they just don't try to apply it or talk to people about how to apply it. But yet the whole philosophy at Harwell was that, and it was there from the word go. It was just built in; you couldn't do anything else.

HAIGH: We'll return to some of those topics a little bit later when I talk about the Harwell Subroutine Library and its relationship to the users. Let's just get some big-picture information on your return to Harwell. So you came back as something called a Senior Scientific Officer.

DUFF: Yes, that's pretty good. If you're recruited as a fresh Ph.D. into the civil service in Britain, you're a Higher Scientific Officer, which is a lower paid position. So I guess I had been out just about long enough. Then it was very age-related and experience-related as to what grade you came in at. So I think there would have been a bit of a debate, I'm sure, about whether I would have been somewhere up from the bottom of higher, maybe, but "higher" or "senior". I can't remember if I did any bargaining there particularly, but for one reason or another, it was decided that it was appropriate to make me Senior, probably quite low down on the ladder, though.

HAIGH: So were those kinds of titles just important for how much you got paid, or was there a status system within the culture?

DUFF: Not really status, no. Pay, yes. It was definitely pay. I mean, of course then, my first paycheck was something like £2,000 pounds a year. It was just completely different from nowadays. Then the increase was probably like £200 pounds more for a higher grade. There was an element of status, I suppose, but it was only ranking us against the ladder as a whole. It's not competitive in any other sense. Certainly, I know from managing people that in certain stages of their career, if they're not getting promoted, they'll come to you and ask you why, so that's important. I did quite well to get Senior: not enormously well but quite well. These are scientific civil service grades for the whole of scientific service in the UK. Now, Harwell has never been directly scientific civil service, even in the days when I was there, but it adopted a lot of the rules and regulations from the scientific civil service. NPL was then scientific civil service, but not Harwell, because we were part of a different department; we were part of the Department of Energy.

HAIGH: So your work there, basically, would be doing research and writing papers. Would that be correct?

DUFF: Yes, and giving lectures. We did lecture courses, too. They were quite active for a while. We actually did teaching courses, but not a lot of marking in that. These were intensive courses, three- or four-day courses, to other government scientists, mostly at Harwell but elsewhere as well, who wanted to know about numerical optimization techniques and numerical linear algebra, and that kind of thing. So we had a set of courses we gave, which I think were well received and fairly important.

HAIGH: Was there a formal system of internal consulting to help people with their problems, or was that just managed person-to-person?

DUFF: I think probably then it was person-to-person. There wasn't anything that well organized. Quite often, for example, Mike or Alan, the two most senior people, would have a contact with somebody because they came from theoretical physics, so they would have a chat at lunch about something. Then they would say, "Iain might be a guy to talk to you about that." So then you would get in touch.

HAIGH: You said that early on, the same people might be doing the research and writing the operating systems for the computers.

DUFF: Alan was one of these people, yes.

HAIGH: Was that changing by the late-'70s? Was there more specialization in systems software as a career track?

DUFF: Well, a change in a sense that computer vendors were coming out with operating systems which basically worked. So what was needed from us was maybe modifications to them, but less intensive development. So there was less of that needed.

HAIGH: I imagine there must have been established systems administrators.

DUFF: Oh yes, sorry, but mathematicians were less involved. We were doing mathematics by that time. So people like Alan Curtis, for example, was much less involved in, say, operating system work by the time I got there in 1975.

HAIGH: So the people would still be doing that, but there was just more of an internal separation?

DUFF: Yes, that's right. In the Computer Science and Systems Division that we were in, there were actually computer operators, operating system-type people, and mathematical people.

HAIGH: So before we start talking in more detail about your academic work and the development of the Harwell library, let's just get an overall sense of how your career and the group within Harwell developed over the period from 1975 to 1990.

DUFF: The main dramatic thing was that just about a year after I had came, with Mike Powell telling me how much better I was working at a lab like Harwell than in a university, Mike was persuaded to go to Cambridge as a senior professor by James Lighthill and others. He certainly was very embarrassed about it. He wasn't quite asking me if he was allowed to go, but he really felt quite embarrassed, just because of all the build-up from me coming and the fairly long time it took me to decide to come from Newcastle. And then as I said, I was not there more than a month or so before he starts discussions with Cambridge and then he left about a year later.

John Reid then became leader of the numerical analysis group after Powell left. Alan Curtis by that time was leading the applied maths group, and the two had sort of split, but everything in the office was all mixed up. I think that split was probably a personality thing. Not in a bad sense, but I think Mike needed a group because he was senior; and Alan couldn't be reporting to Mike because he was even more senior, so they had two groups. That was maintained. Things don't change so easily. John Reid, of course, was very much in numerical linear algebra and had been involved in sparse matrices from the early stage, and had been my principle collaborator, and supervisor in addition to Fox in the last year of my thesis. So that worked out very well. I had a very good relationship with John.

The book which was finished in 1986 started not long after I got to Harwell. Well actually, it had started before. Initially it was Reid and Al Erisman doing the book, but I was getting a lot of pressure after my big article in the Proceedings of IEEE which was in 1977—I did that when I was in Newcastle. I got a lot of pressure to write a book from OUP. It was becoming crazy that I should write a book at the same time that John and Al, who were in fact not getting very far with their book. So the three of us then joined together. It was eventually published in '86. It was a long process. But that was partly the thing I was working on with John back in 1976, '77, I guess. I can't remember exactly when I came on board for that book, but it was sometime in the late seventies.

HAIGH: I'll ask you more about the book later. In 1979, you got another bump to Principal Scientific Officer.

DUFF: Yes, that was a two-stage thing. So one interview I failed in a sense. They had really quite severe interviews by many physicists, not people in your subject. I can't remember; it was sort of a half failure. It was more or less, "This guy's a bit young; send him back next year" sort of idea. But the head of the department appointed me as a Temporary Principal Scientific Officer from the moment that happened. So I was then paid in that band. I think it was called temporary PSO, actually. It was confirmed the following year by another board. So that actually was a big shock to my system, and ever since then I've been nervous about boards where I wasn't before. I must check who was on that board and see that they're all in their graves. I was really upset.

HAIGH: So in general, then, you were having an unusually rapid rise up through these various grades?

DUFF: Fairly rapid. I guess it was normal for people were maybe "destined for great things", you could call it. But then the organization was destined for great things. It wasn't absolutely meteoric—I didn't become Head of Lab or anything, but it was pretty fast.

HAIGH: Then in 1986, you were Group Leader of the numerical analysis group.

DUFF: John was well off retirement, he retired not that long ago, but I don't think that in some sense he fully enjoyed being group leader, to be honest. I was, as you know, very political, and that included political in the workplace, and I was quite interested to take over the group, so John was really happy for me to do so. Nothing much happened. He became essentially emeritus, even though he was well off retirement, and I became group leader. I'm trying to remember when I became Senior Staff banded, which is the grade above Principle Scientific Officer. I guess it was about the same time. I guess it was about '86, was it. I think it was a coincidence.

HAIGH: Yes, 1986.

DUFF: That wasn't related to being group leader. I mean, the two were not associated in the sense that one led to the other or vice versa, it was just that sort of time. So that was pretty good. That, I did get on the first interview. I came out of the interview absolutely convinced I had failed it because of my previous experience, and there was no trouble. So I got that. The group leader thing was really a discussion between John Reid and myself. We felt that I could bring “new blood” in as group leader, and he was really happy not to be group leader, so we came to that agreement. It worked out well. It was always the issue about what would happen with my senior. I mean, John is nearly ten years older than me because he had done national service, so he missed out on some work years for that. It was always going to be interesting how it would work, me being his leader for a change. I would like to think it worked very well and I would think he would agree as well. So I think that went okay. And one thing I did when I was group leader was I every much emphasized visitor programs. Any money I had I used to help support visitors that came in for a short while. That worked out very well. Well, I think it did anyway.

HAIGH: How did the general mood at Harwell change over that period? I imagine in 1975 people had been full of optimism about the immediate future of nuclear power.

DUFF: Oh yes.

HAIGH: And clearly, that must have begun to change by the 1980s?

DUFF: Yes. So the lab was diversifying scientifically, and they had such a wonderful group of people, as I keep mentioning, that that wasn’t a hard thing to do. A lot of the workers were fairly theoretical, and their connection with nuclear power was maybe not immediate anyway, even back in the old days. So that continued very much, and in 1986, I would say some groups in Harwell were beginning to feel the pinch. So there were some issues there. But not our group, thanks to Eric Taylor, our division head. He was not a numerical analyst, I don’t think he was even a mathematician; he was a nuclear physicist, I think. Department heads were all-powerful in these days, as the civil service works that way (even though Harwell wasn’t technically that). He was very well disposed towards us all, towards Alan Curtis, to John Reid, to myself. And in fact, I was able to recruit two new people in that period: Nick Gould and Jennifer Scott. So we had built up the research team very well. They’re still with me to this day.

HAIGH: So on that topic, how large was the numerical analysis group?

DUFF: Well, that was the one reason why I was a little obsessed, a strong word, but why I was quite keen, to become group leader, because John wasn’t interested in recruiting new people. It just wasn’t his thing. So we had these two people who were working primarily with the Harwell Subroutine library, who were Mike Hopper, who also did some typesetting systems, and Sid Marlow, who did the day-to-day libraries. They were always in the group. John and me were there. Pat Gaffney was there for a while. That was the person who was the student of Mike Powell, but then he left. But it was very small. That was the trouble. So I was really keen to get more people. We had had visitors, though it wasn’t that I initiated getting visitors, I just emphasized it more. So we really needed to get more people, with the help of Eric Taylor. Obviously personality-wise we got on with him, and he was, I’m sure, bending rules and making us protected in a way that Alan Curtis’ group was not. Because about the time I took over as head of numerical analysis, a fellow who was about a year younger than me called Ian Jones took over from Alan Curtis as head. Alan was continuing, but he took over as Head of the Applied Math Group, as it was called. Their group definitely had more commercial pressures almost from when he took over. But we were immune, and Ian was actually a bit jealous of us, in the nicest way. I mean, he would be openly jealous, let’s say.

We were very protected until Eric retired, which was 1989, I think. And then the writing was on the wall! We were being told in no uncertain manner that we had to earn our keep and that researching papers was not enough and not wanted. It was really in these words! I mean, it was absolutely strong. And I had just recruited these two people not long ago, around 1985. Well, actually I think I recruited Nick before I was group leader. I had just taken up this commitment with Strathclyde as a visiting professor and at CERFACS as a group leader to try and keep my research interests going, which was actually done with the blessings of Eric Taylor before he left. I would never have got that through afterwards. So everything was also set up personally for me quite well, and for the group, until the point at which Eric retired.

Mike Norgett took over, Mike was a theoretical physicist and personally a nice guy, but he was in the new mold— he had been educated that research wasn't enough. Sorry, he had been re-educated because I guess he started in the old days. He seemed to believe that, and wanted us to earn our keep. That was when I started going to Hong Kong. I was in Hong Kong six times in three months, sometimes on a day trip, if you can believe that! 17-hour flight there and the 17-hour flight back, and only 12 hours in Hong Kong itself, that type of thing, working for a company called BIS Macintosh which is a famous group of management consultants. That made me realized why I did not want to work for them. Then I was a technical consultant for supercomputing.

HAIGH: Maybe we should come back and do 1989 into the '90s later on then.

DUFF: I was into supercomputing by that time, which was the big thing. We had got a big Cray at that Harwell, one of those huge Cray 2s, a big liquid-cooled thing. So I was seconded to BIS Macintosh working for the government of Hong Kong evaluating supercomputers. I was flying off to Hong Kong all the time. In fact, I actually left for Rutherford Appleton Lab later than the rest of the team because I had a commitment to fulfill the contract with the government of Hong Kong before they would let me away. So I went out in June and they were there in April.

But let's complete the Harwell thing. The move to Rutherford was really very fraught. I was at a fireworks party in this village with a guy who was pretty senior at Rutherford. Now, he's head of Isis and Deputy Director at Rutherford, Andrew Taylor. He was living in our village, and he knew me quite well. He said, "You're the kind of person we need at Rutherford Lab." I said, "Yes, well maybe this is the place I should be going because you still do research there." He said, "Oh yes, that's why we want you." So that started me off talking to Brian Davies, who was Head of Computing, effectively, and head of a big Cray 1 at Rutherford Lab. I was then in protracted negotiations with them and Paul Williams, who was then director at Rutherford, about moving our whole group over. Our whole group, then, ended up being about seven people actually. That was myself, John, Nick, Jennifer, Seamus Considine, and the two who were in some sense more junior people, even if they were older than me, Sid Marlow and Mike Hopper. At one time, we thought we might get everybody over, but as luck would have it, they had a funding crisis at Rutherford for the first time in history. We squeaked in by the skin of our teeth eventually. We couldn't take everybody; just the research people. We had to go to the SERC—or was it the SRC at that time—council, debate at the highest levels, we had to write to MPs, we had to do everything to get our visibility level up enough for us to be accepted and we got transferred. That was a big effort. What we were saying is if we remain at Harwell, then our research will not be able to be done, so if you want this research done for the benefit of Great Britain, then we have to move. They bought it enough and moved us.

HAIGH: And after the move, were things much as they had been before?

DUFF: Much better. I mean, far better, because immediately the pressure to earn money and do commercial work was removed. We were there to do research again. For the year and half, maybe a year since Eric Taylor had retired, we were under constant pressure. That's why I had to do the stuff for BIS Macintosh. It was quite entertaining. I didn't mind flying off to Hong Kong, jet-setting around and that, but I wouldn't want to do it for very long. It was quite long enough. It was really a six-month period I think. I was sort of flying out on a weekly basis.

The pressure was on the other people as well. Nick had to go to a paint factory in Slough and watch paint drying, and model it or something, to get money. It was more or less that we were forced to earn by commercial contracts our salary plus. That was the way it was going, the way it's continued to go there as well. They're not interested in the science; the science can be absolute crap. It's just money they wanted. The thing that was really galling is that at that time it was scientists managing us. Since then, they've brought in people who know nothing about science, which is even more iniquitous. But at that time it was actually scientists, like Mike Norgett, who were in control, and they saw the writing on the wall, I guess. I just don't know. It was sickening that nobody could fight it, and I was too junior at that time, even as group leader, to fight the whole trend. It just bewilders me, but it's happening in other places. The government, not the government necessarily, but local people, just destroyed something that was excellent. Nearly the whole of theoretical physics left, and they founded departments in universities which are the best in the world now. University College London, for example, was almost completely staffed by ex-people from Harwell.

HAIGH: So was the impetus towards that commercialization coming from the government itself?

DUFF: Well, it's not entirely clear to me. I used to blame Thatcher. I guess there was some influence there because they wanted to move people out of the public sector. We were definitely public sector, even though we weren't civil service. Nuclear power was on the wane, but really since then, it sort of fell apart. But anyway, it was starting not to be the main activity at Harwell. There was a lot of other energy research. I'm not party to what was happening at the higher levels, but certainly our management was implementing it as if they believed in it, which was a sad thing. So they never said it to us as a diplomatic, "Because of the government we're going to have to do this, sorry." They moved into it as if it were the right thing to do when it completely wasn't. Everybody could see that, anybody with half an ounce of scientific sense.

I should say that, I'm really not bitter, of course, personally. I wasn't because we got to move. But I think that in terms of the destruction of a scientific environment, this was tantamount to destroying the Buddhist statues in Afghanistan. It was really shocking, shocking.

HAIGH: Once the group moved to Rutherford Appleton, has it stayed basically the same format that it had been in Harwell prior to this push for commercialization, that you were judged on academic merit?

DUFF: Yes, I would say so. Yes, largely. The library has sort of evolved somewhat since the early days, but the actual work we were doing was as close as to what we had in the early to mid-'80s at Harwell, yes.

HAIGH: Now we've got that basic framework in place, then, let's move and talk in more detail about different aspects of your work at this period. Let's start with the library itself.

DUFF: Yes. It's now called HSL because, for obvious reasons, calling it Harwell Subroutine Library, which it used to be, is a bit silly now because we've been away from Harwell for so long. So the name of it now is the HSL. But it was the Harwell Subroutine Library from the early days. It was maintained, distributed, and supported as a general mathematical software library eight years before NAG was even dreamt of. Brian obviously has to admit it when you face him with it, but he would never tell you about it. In fact, it was about the same time as the Bell Laboratories library was founded, about 1964. So I think HSL is actually probably the longest standing library project that's still continuing as a library project. It was, of course, originally developed by Mike Powell, and to some extent Mike Hopper. Hopper implemented the structure of the library, and Powell was, let's say, the research brains behind it and supplied a lot of the early software. Alan Curtis was very involved and, of course John was, too. But it was Mike Powell who was in charge of it as group leader of CSS at the time. He was section leader at that time in 1963.

[Tape 2, Side B]

DUFF: The people doing computational science at Harwell had needed to do numerical computation, of course, and at that time, there were no libraries. NAG wasn't even thought of, as I said, and IMSL wasn't around. There were no other mathematical software libraries, so people wrote their own codes. These people in the main were physicists and very intelligent, as I said, but not specialists in writing numerical code. There were ten codes around, say, for solving a linear system (not a sparse one even, just a dense one) a central problem that Department of Energy labs addressed through EISPACK and other projects in America. So what Mike established was the fact that nobody else at Harwell should write a code. It should almost be forbidden, and all the codes should be written and supported by the numerical analysis group. He was largely successful in that. The codes worked, Mike was one of the founders of modern-day numerical optimization, him and Fletcher. They had brilliant algorithms that were quite unlike anything that had ever been developed before, new algorithms which they coded. They coded well, though Mike's coding was a source of great amusement to everybody because nobody could understand his codes. They weren't exactly transparent, but they worked. He had some really very good high-quality work, completely world class. That's why he got his FRS. The quality of the software was above any existing in America at that time. As well as numerical optimization, the original sparse matrix work of John and Alan was included the library, and that was also world-leading at the time. So it got a very good name for being one of the only mathematical software libraries, and one of the best in its fields. That was way back in 1963. I was asking Mike recently for another purpose, and in fact, there were no external users of the library until '64, when they started to ship it elsewhere. In fact, it was more freely shipped then than it is now. It was not at all commercial, and now it's a bit quasi-commercial.

HAIGH: So in these early days, what machine as being used at Harwell?

DUFF: The original HSL was in what you might call IBM FORTRAN, it was specially designed to be efficient on IBM equipment.. It's still in FORTRAN now but it can run on other machines with FORTRAN compilers.

HAIGH: So in those days, do you think it could have been a 7090?

DUFF: Yes, we've had that. We had a 7090 at one time, and the 370 series, as I mentioned as well.

HAIGH: 1963 and '64 are a little before the model 360s could have been in place.

DUFF: Oh sorry, that's right. 7090 was right. You are correct.

HAIGH: Do you know if in the 1960s any efforts were made to port it to perform well on other platforms?

DUFF: Very little, I would say. I think there were some efforts made to port it to other platforms. In the later 1960s, by the time I was a student there, there was more of an interest that it would run on other platforms because people were using other platforms. But in terms of trying to make it efficient on other platforms, I don't think much was done. There was more happening when I arrived in 1975. It wasn't me that brought it in, but people were getting more concerned then.

HAIGH: Did you have any experience with the IBM SHARE library, which would also have been for these big machines?

DUFF: Yes, we did. Bob McLatchie was chairman of SEAS (as SHARE Europe was then known) for a while and he was in our department at Harwell. For a short time around 1989 he was in charge of the Harwell Subroutine Library when they were trying to commercialize it, rather unsuccessfully, as Mike could have told them. So in the early days when the HSL started, I'm not sure there was very much in the way of software of any interest from that. I don't know when the SHARE thing started; you would know better than me.

HAIGH: Very early. It started with the IBM 704, and then it really got up steam with the 709. So it was firmly up and running in the latter part of the 1950s.

DUFF: So I don't think that that was particularly good for the sort of what we would call then "high performance" computing, which was the need at Harwell. I don't think the SHARE library addressed that very much. But we certainly were aware of it, and people used it.

HAIGH: One of the problems that I've come across in other context with it was that the code was contributed by all the different IBM installations, so some of it would be very high-quality and a lot of it would be terrible, and the effort involved in finding out what was very good was not necessarily worthwhile if you could be writing something yourself.

DUFF: Well, I think the strength of HSL has been the quality control—we haven't really made it fantastically automatic to this day in a sense, although we do put it through various tests. But in the early days, it was even less automatized, if you want to call it that, and it was all on the strength of high-quality research, high-quality, intelligent people like Mike Powell. There really was an effort to make it high quality, and there always has been a big effort for that. Obviously we do find bugs—you can't avoid that in big codes—but very few and we solve them pretty quickly if we hear about them. And we do maintain and support the codes, and always have done, and still do. That's really what the strength has been and why it was adopted at Harwell. I mean, we're talking about physicists who were really quite capable of writing their own codes, but they would rather do physics, and if something was wrong with our codes, they would have immediately abandoned them and done their own thing.

HAIGH: Some of the people I've talked to who were responsible for library efforts in other places, for example, in some of the US national labs, did speak sometimes about difficulties in persuading the users to utilize library routines rather than writing their own, and they talked about things like training days and newsletters and consistent documentation standards. So in this case, were the users invariably happy to use a supplied routine, or did you sometimes have to sell it to them?

DUFF: Because all that came as a package, there was initial reluctance to use it—anything that comes in a black box. But that was overcome to a large extent, and I think that was really the efforts of people like Mike Powell and Alan Curtis in relating to the physicists. Because of the quality (and people began to recognize that soon), the robustness, the fact that we were willing to not only resolve any kind of bugs or make any additions to the code to help them, but we would also discuss sometimes the computational problem and the formulation of it and would be working with them.

HAIGH: So you think the quality was such as to really sell itself?

DUFF: Yes, and I think it did.

HAIGH: So imagine you were a physicist and you wanted to use some of the functionality from the library in, say, the early 1970s. What would you do? Would you go along to the computer center and pick up a deck of cards with the library routine on it and pick out a particular code to stack with your own source file in the hopper?

DUFF: I think this Harwell User's Workshop had a system for allowing there to be essentially compiled code you could link into, even from an early stage. So I think we had compiled libraries from a very early stage.

HAIGH: Those would be resident on a disk so you could just insert a call into your own code?

DUFF: I believe, essentially so. Or there's another possibility that these were held in the computer; people didn't have to come to us for decks of cards. So whether they were held in compiled code, and I think they probably were from an early stage, because don't forget; we tended to be just using the one machine most of the time, the big IBM. So we had a little problem every time we got a new IBM, as we had to compile the FORTRAN and make sure it worked again. But I believe essentially it was compiled modules from whenever that started and computing was available, I'm sure we were using it, but I'm not 100% sure. At the very worst, the FORTRAN would be available on a disk; the person would take the FORTRAN code off the disk and compile it with their program. I think that probably did happen a bit as well for a while.

HAIGH: Were there, say, printed manuals in ring-bound folders in the computer lab that you would go and look at?

DUFF: Yes. Absolutely. They might even have some still.

HAIGH: Did people put a lot of effort into documentation? Were there any standards?

DUFF: Yes. We did from an early stage. I just mentioned Mike Hopper, who is in his late 60s now, maybe even 70, and he still comes along as a consultant for us for the library. He manages the library on a day-to-day sense, and also manages Sid Marlow who has since retired and is not around anymore, although he lives not so far from here. So Mike and Sid together kept the library going at that stage. In fact, that's what John Reid's doing now. He's employed by us as a consultant to do that. So it's gone full circle in that sense. Mike Hopper became very interested in typesetting. This causes great laughter by friends in America like Gene Golub and Jack Dongarra, but we actually had at Harwell at that time before LaTeX a typesetting system that was vastly superior to LaTeX. TSSD had macros and everything. We typeset our documentation in this. The dates again are difficult, but certainly it was around 1980, well before LaTeX was developed. In fact, I think we had some possibilities to market it. John was the group leader at the time, and I wasn't involved in any of the discussions, but I feel that we could have marketed that, in the sense of distributing it. I'm not sure we could have beaten Knuth at his game because he's a rather powerful figure, but we could have had a very good typesetting system. It still exists! We still use it to this day.

HAIGH: There were also some popular Unix tools coming along, I think in the mid-'70s, such as TROFF and NROFF.

DUFF: Well, they're not very powerful, but yes. Were not at all a Unix lab until about 1989. Just before we left Harwell, we got a donation of six SUNs from AMACO, an oil company who had obviously some money that they didn't want to be taxed on or something and they were giving it away to various places. John Reid knew Jerry Wagner through his embryonic FORTRAN work, I think he was called from AMACO, and he chatted him up. We got six SUNs as workstations, which Harwell at one time wanted to stop us taking them out. We didn't quite smuggle them out the gate. We didn't really use Unix. The closest we got to them was PDP machines, so that's another story that I've got.

Back in the early '80s, I was the only person to use email, and we used a PDP machine—PDP 10 or 11 or something—over in another building. That used the VMS operating system. That was another flavor of Unix, I guess, and that was the closest we got to Unix before we actually got our Suns, but since then we have used nothing else.

HAIGH: So did this typesetting system have a name?

DUFF: Yes. TSSD: Typesetting for Scientific Documents.

HAIGH: So that's the documentation. Now, how about the sources of the routines? So you've mentioned the names of a number of contributors and library people. Were all the routines in the library coming from this small group of people involved with the computer center?

DUFF: More or less. Some of them were collaborative research with people outside, and the people outside might have contributed also to the software. But there was certainly no third-party stuff coming in at that time as far as I know. We've still not really had much totally third-party stuff, although some of the more recent codes have been almost entirely written by people who were not at Rutherford, although in conjunction with people at Rutherford. So now it's a little bit more than just in-house, but it tends to be only connected with our research.

HAIGH: And were there any cases where people elsewhere in Harwell, in divisions doing different kinds of science, might contribute?

DUFF: That's right, that did happen. There were quite a few codes that a theoretical physicist, who was in the same department with us at one stage, of course, would have developed, and then we would have done some quality control on it. So there were people outside, but from within Harwell who were involved. And I would say that probably we're moving more to being interested in third-party codes. So for example, we had someone from my team at CERFACS visiting us. He's no longer at CERFACS, he's at the local university, but came from Toulouse, and we paid him as a summer student. He's a *maître de conférence*, a junior professor over in France. So he came over, and he developed a scaling routine for the library by himself. That wasn't a giant research project; we were just interested in having someone who did it. So we are moving a little bit in this direction.

HAIGH: So in the 1960s presumably this Harwell Software Library would have been the only library available to users within Harwell?

DUFF: Yes, the only library available to users anywhere, to some extent. Although I think Bell Labs had the PORT library.

HAIGH: Presumably, yes. I think to some extent in those days every computer installation had its own home-grown routines.

DUFF: Yes, there was a bit of that, sure.

HAIGH: Now later on, when high-quality codes started arriving from the outside world like LINPACK.

DUFF: Quite a lot later, that was, but yes.

HAIGH: And somewhat earlier for EISPACK.

DUFF: EISPACK was earlier, yes.

HAIGH: When software from the outside world that was worth having came along, would you just say to the users, "All right, here on one side is the Harwell Library, here on the other side is EISPACK," or would you make an effort to try and present it as being part of one unified library, to pull the documentation together, that kind of thing?

DUFF: We sort of tried a bit. Alan Curtis took some of the LINPACK codes with appropriate acknowledgements and created a Harwell interface or HSL interface to them and we distributed that. We also encouraged people to use these libraries and had them on our machines, in particular, EISPACK and LINPACK, the two you mentioned. But EISPACK was in the mid-1970s and LINPACK was '79, so they were certainly a long way after the HSL started.

HAIGH: Yes, so for a long time then, this would have been just the only library that people were given access to within Harwell and they…

DUFF: Unless there was the odd stuff in SHARE as you mentioned or somewhere else.

HAIGH: Now, another interesting feature of the libraries has just been the organization. It seems that different libraries have been divided into chapters, or on some system that has been inspired by the Dewey Decimal System to try and bring the stuff into some kind of sensible taxonomy.

DUFF: Right. So that was done right at the very beginning by Powell. The reason was that what you didn't want was to have a routine called "solve", because users in their own code could have half a dozen routines called "solve", and then there could be all sorts of confusion. So right from the word go we started with chapters in the library. Interestingly enough, we've got a different system than NAG. We were constrained by six characters for FORTRAN names at the time, so the first two identified the library chapter. So MA was "matrices" (some of them were mnemonic and some weren't); VA was unconstrained optimization or something; EA was eigenvalue problems. So we had these libraries set up more or less unchanged from when they were initially set up. We have introduced a couple of new categories because of new routines. So there were two letters followed by two numbers running up to 99. We haven't quite exhausted any of the categories yet—we're getting closer; we're up in the 70s now. So that was two letters, two characters, and then the entry call, because there might be several entries—A, B, C, D, a letter—and then the last letter, which was a bit extravagant, profligate really, was to have a D if the code was double-precision and nothing, a blank, five characters just if it was just single-precision. So we always carried the double precision 64-bit and the single-precision 32-bit. So that was how it was named, and that avoided confusion. It also avoids confusion with NAG as well which is quite nice, because they've got a letter followed by two digits followed by… anyway they've got a different way of doing it. I think Mike probably got his naming convention from somewhere, but I don't recollect.

HAIGH: You mentioned Reid and Powell several times. Could you talk a little bit about them in terms of personality, interests, those kinds of things?

DUFF: Well, Mike is a very dominant character in optimization and was only the second fellow of the Royal Society in numerical analysis to be appointed, the first one after Jim Wilkinson. He's one of the founding fathers of the sort of numerical optimization that we know today. Mostly, of course, his work was on small-scale problems by today's standards, and even today his strength is probably in analysis and working with small problems, although he has since acknowledged that you can get large ones. He's a very keen golfer. He is now an Emeritus at Cambridge and retired, he still does some math but spends more time playing golf. He's the sort of person that you can get to like, and he's a very friendly, family oriented, nice guy. But work-wise, he's as sharp as blazes, and if you get on the wrong side of him work-wise you'll be cut to shreds, rather like Alan Curtis used to do before. So he can be intimidating. He was on one of my later boards and I was nervous as hell, but it worked out fine, in fact. I know that Nick Gould absolutely doesn't want him on his board because he's over-nervous about him and he's in the same field; he's in optimization. So Mike's quite the character that way, really. He's certainly a nice fellow, though.

I've got to know John very well. We weren't very close really when I was a student, even when I was in Newcastle when I came back on visits to Harwell. He was my group leader for a while, then I became his group leader. Then when he retired, I recruited him on a part-time basis as a consultant for the library other research activities. We're still trying to do the second edition of our book with Al, so it's a bone of contention all around. It's not especially between him and me; it's sort of me and me, and he and he, and everything else. I've known him and his family for many years. One of the pictures I was looking through is his daughter as a newborn, who is now in her 30s. So we've known each other very much throughout the years, and I would say we are friends as well. He's very English, and went to public school, although I don't think he was happy being in public school. He did national service; he was just old enough for that. He was, I think, nine years older than me. But we do get on well. He called last week and a couple weeks ago when I first came back from the hospital to check if everything was all right. I don't know what else you want me to say about him. Again, his main area is numerical linear algebra. A thing that I've not been entirely happy with as group leader and even since, is that he got completely sucked into the FORTRAN language development and he was, in fact, secretary of the X3J3 Committee in the United States; he was chairman or secretary of the ISO Committee and he's been very heavily involved in FORTRAN 90, FORTRAN 95, FORTRAN 2000, and he writes this book with Metcalfe. So his life is half FORTRAN and half consulting for us. We don't pay him to do FORTRAN. So he's very much involved in that, and that's to the detriment, to be honest, and he would admit to it too, of his sparse matrix work. He still does some sparse matrixing, but he's got very involved in FORTRAN.

HAIGH: So back to the library. We talked about documentation; we talked about the sources of the routines. Could we talk a bit about the procedures that you'll go through to add a new routine, debug it, test it, or to update an existing routine?

DUFF: Well, these have changed remarkably over the years, as you can imagine. We have the procedure and we have the code. Additionally, we have to write a small spec sheet test, and the code would actually go into the specification. So the documentation will always have a specification sheet which could be from one page to 20 long, describing the code, the parameters, everything else, and then there's a shorter catalogue entry. In the spec sheet, there will be a small example code that somebody could run just to check that the things compile. It's not exhaustive at all, it just runs the code. Sometimes we have two spec sheet tests if there are completely different functionalities that are important and different. Then we have so-called exhaustive testing, which you have to do, which can be quite horrendously complicated. The idea is to execute every statement, not necessarily every path, though we would like to do that. Or if there is a statement that is not executed for some reason, to be able to look at it and say why not. But normally we test them all including even the complicated error returns. That actually doesn't take longer than designing, thinking of, and writing the code. But it could take almost longer than coding the code, if you like.

HAIGH: Did you produce any special tools to accomplish that testing?

DUFF: We have, yes, over the years. Some of the lab routines are profilers that will take the FORTRAN code, annotate it, and then print out a file telling you how many times it's executed. So just adds, counts every statement and branch point. It's automatic. I think John and Alan wrote it. But anyway, it was slightly before my time. It was about the time I was there, but I wasn't involved in working on it.

Now, there's other profiling software that's available from vendors of FORTRAN compilers, so we'll use that. One thing we used in the early days when we were writing FORTRAN 66, as it were, was PFORT. That came from Bell Labs. Now the people who actually market the library for us, which is Hypertech, use CVS or something for quality control, for keeping up and updating the library.

But as I said, every time we do work on the library, and Mike Hopper helps us there, and we tighten things up a little bit trying to do it more rigorously. We have done software for NAG that's been put into NAG, and so we've gone through their procedures as well, and we've found them to be rather similar to ours.

HAIGH: Once a routine was in the library, would it be common for changes and improvements to be made to it?

DUFF: In the old days (the old days being until a year ago, about), no changes could be made to a routine that influenced the interface. Obviously any bugs inside could be fixed; there could be a little bit more internal stuff, maybe a slight twitching of the algorithms; internal constants might be changed slightly, so it could cause different output. But basically the appearance had to be quite unchanged. And if we made any change to the interface, we had to introduce a new number. So for example, we used to use common blocks a lot, and then we decided that we should make the library thread safe, so we had to do away with common blocks. Any time that we took a routine and made an identical routine by passing variables as parameters rather than through the common block, we changed the name. So we went from EA12 to EA22, for example. We usually kept some sort of mnemonic for our sanity in the name change, but we changed the name.

So nowadays, we've got a new idea that we've started last year, which was present in other libraries before but we haven't done, and that's to have an explicit release number. There's three digits to that. There's one that's just minor changes, almost changes to comments or something, which is the third field. The middle field would be localized bug fixes and various things of that kind. The other one could be that there is an addition to the interface. We try and make it upward compatible all the time, otherwise we change the name, but we would allow there to be additional facilities, maybe a new entry point or something else. That would be a first version change. So some of the codes are in version 3.0.1 or 3.1.0. It's because there's name recognition. There's a code that I developed a little while back called MA57 for symmetric and indefinite matrices. It is used all over the world by people in optimization, and they know of it as MA57, and if I tried to say, "Well, I've got some other name," they would be confused. So we've kept the names, but put in versions.

HAIGH: So prior to that, the new version would have a new number, so that if you wrote code that relied on the old version, it would still be there.

DUFF: Yes, we kept them. We never threw away the codes.

HAIGH: And then if you wanted to take advantage of the new version, you would have to rewrite your code to the new one.

DUFF: Yes. So what we've done now is that the old codes are still around, but we've put them into an archive library that's independent. So if somebody really wanted to use one, they could go to the archive, which is totally free for any to access.

HAIGH: Was that because over time the library had just been building up with just too many routines?

DUFF: Partly. It was getting really so large to maintain, it wasn't really even maintainable. Every time we brought out a new release and had some new feature or something new, like when we went thread safe, it was just horrendous when you've got 400 routines or something; it just takes forever to change them. Even if you've got automated tools to help, it still takes a long time. Off the top of my head, I can't remember how many routines we've got, but let's say it's down to 100 rather than 400. Certainly, we're trying to reduce it.

HAIGH: So that brings us to the broader topic of the changes to the library over the years and the ways it's evolved. Did it expand appreciably in terms of the general coverage and different kinds of mathematical routines?

DUFF: It evolved rather than anything dramatic because we started off with having software, for example, for small dense systems and little bit for sparse systems. Then the sparse codes developed, and now we've got around ten routines for solving a sparse system, depending on the nature of the sparse system, if it's coming from a particular application, a particular way the matrix is stored, whether the matrices are definite, indefinite, asymmetric. So we've got a whole number of routines for that. We had an expert in eigenvalues working for us for a few years at Rutherford. When he was there, we would put in some quite new and innovative eigenvalue routines. He's gone now. And so we'll get some new routines and new methods coming in. So it kind of evolves rather than any radical change.

But the main theme all along, since we started in 1963, was to support scientific computation, simulation, and large-scale—of its day—calculation. The large-scale wasn't very large in '63. So that theme has continued; we have kept that going. So it's been quite different from NAG in that sense. NAG is really a general-purpose library trying to be all things to all people and essentially cover all areas of mathematical software.

HAIGH: So in terms of comparisons, are there areas where you would say that traditionally, the Harwell Library has been stronger than NAG?

DUFF: Oh yes, very much. Optimization, of course, would have to be one from the early days of Mike Powell. Sparse matrix work is definitely another one, especially direct methods. We're probably equal in iterative methods; they've got some iterative codes as well.

HAIGH: So do you know at what point NAG or other major libraries would have started to include sparse matrix codes?

DUFF: Yes, well the first sparse matrix code that NAG ever had was one that we called MA18 in our library, which-- oh my gosh, this is embarrassing! I can't remember if MA18 was one I co-authored or whether that was the Curtis-Reid one before I got there!

HAIGH: I have your list of publications. Would that help?

DUFF: Maybe, maybe not. No, it was MA18. I think that was a Curtis-Reid one. So that was a very old routine. That one is in NAG, plus some other ones. At that time, we had discovered that the quality-control that NAG had was a bit tighter and so we adapted their approach for the whole library. I'm trying to remember when that was. This was a long time ago. I mean, we're talking here, 1980, before I was Group Leader, I'm pretty sure. Early '80s, let's say. So they incorporated it fully into the NAG library. We were a bit miffed because, well, first of all, we were new and advanced, but we got no feedback because NAG doesn't work that way (or didn't then, anyway). We thought we might get some fame and fortune, but we didn't.

HAIGH: So you think that…

DUFF: Well, it just became a NAG routine.

HAIGH: … you weren't acknowledged properly?

DUFF: Well, it was probably in the documentation, but very little. So it was known as a NAG routine and not a routine from the Harwell stable. So the second time around, when NAG came back to us to ask for an update to it, we then negotiated and changed our collaboration. We developed a library called HSML which is a NAG-distributed library of our codes, and that included MA28 (an update of MA18, that was one I had been writing) and some more of our newer routines. That sort of worked. They did distribute it. It was quite clear in the documentation that it was ours this time and it was identified separately from the NAG library. And we sold a few and we made some money from it. This has now sort of fallen apart a little. I mean, NAG is now looking elsewhere. There are other sources of sparse/direct codes existing now. Whether they're better… I would maybe argue not. But there are people in America who have now come into this side of sparse matrix work and developed codes. So NAG can actually look around and get sparse codes elsewhere. So the HSML is not really appropriate anymore, and I think it's totally dead now.

HAIGH: Did any of the Harwell Subroutine Library codes make their way into the IMSL library in America?

DUFF: I think so. We had a lot of inquiries. Dick Hanson, who was involved with IMSL, was a very good friend of John Reid's, and a friend of mine as well. We had long talks with him about codes, and that was all to do about commercial arrangements. I think we have some codes there. But it's not something that's been happening particularly recently.

HAIGH: So you had mentioned the specific products distributed by NAG on a commercial basis. Was that the only commercial distribution that the HSL received, or were there other mechanisms?

DUFF: No, that was the point, and that was why it was all a bit delicate, really, because since the days when Harwell was in its death throes before we left, they were trying to get money out of anything—they would get money out of their grandmother if they could have sold her. I mean, really. It was as bad as that. They tried very hard to make money at the library in the late '80s when we were still with Harwell. We didn't get a lot of money for trying to sell it. We maybe got a few customers, but you couldn't sell it for a lot of money. It's still the case to this day. As I mentioned earlier it's remarkable that NAG's selling enough software, I would think, but they get more money from compilers than from the main library, I believe.

HAIGH: Were they trying to sell it directly to users, or was the idea that it would be sold to developers who would then compile some of the code into specific applications?

DUFF: Sort of both. The embedded software aspect was where the money was and still is. So when we left to come to Rutherford, Harwell claimed that HSL was their intellectual property and, to some extent, I guess that was right since we had been working for them. But we wanted to continue developing the library and use that as a basis, and they wanted us to continue to supporting it, because they had many users at Harwell that used it still, in spite of the fact that they're doing less research. So we had to get a compromise, and we managed to get one which sort of by and large has kind of lasted. The idea was that Harwell would continue to have marketing rights and we would have the rights to use it and continue to develop it, and we can use them with all our internal people and research centers in the UK universities. I can't remember if they had the sole marketing rights; I don't remember the exact details these days of the original contract. That has developed a bit since, but in general terms, we do have the use of it.

So Harwell is going through its death throes, there are departments closing, opening, people are being sacked. All the people that ever managed the library kept getting sacked or moved and the library got passed to somebody else. They tried to sell it internally as a great asset, and it wasn't really. Until eventually about five years ago—I think—a group took over the library. It used to be part of the heat transfer and fluid flow service and then became a sort of quasi-company called Hyprotech and then a real company called Hyprotech, which was later taken over by Aspentech. We've had a very good relationship with them for five or six years. They use it themselves in the chemical process engineering work that they do, which is a big area for direct sparse codes. So they use quite a few of our codes heavily, and that's included in the deal—they can do that for nothing. But they also market it, and they're marketing their own codes and big codes in process engineering, so this is like a sideline. I'm not sure they're very active marketers, but we get plenty of interests thrown up. Not big money—we're talking less than a 100,000 a year—but the money we make is usually from embedded sales. Somebody either wants to build a product for their own company or even build a product to resell and they want to use good robust numerical driver software inside it. That is where the main money comes from. But some people buy the whole library, and that's nice, so we don't charge so very much just for their own internal use.

HAIGH: And in the cases where the code was sold, was it being sold with a license that let you use it indefinitely, or was it renewed every year?

DUFF: Yes, the way we did it was just a perpetual license, a single payment. No hassle with that. And no guarantees for upgrades. I mean, there was some amount of support offered, and any bug fixes that came would be offered, but if we developed a package, let's say, and made it into something else, they didn't get that as part of the deal; they would have to buy a new one.

HAIGH: So then moving back in time a bit from these efforts to commercialize the library, how about use outside the lab on a non-commercial basis? My impression is that for a while the library was fairly easy to get for non-commercial purposes.

DUFF: Well, initially it was totally easy to get. Mike Powell just sent it out back in the 1960s, and even beyond that. Then from the time that Harwell was in its death throes, it was really difficult to get it to anybody, because we didn't know who was in charge. It's improved a bit now. Not quite back to the old days, but now our group is supported almost entirely on EPSRC grants, and partly as a result of that, and this was in agreement with the people at Aspentech/Hyprotech, we make our library completely free and available to any .ac.uk address. It's monitored access, but they can access and download any codes from the library for their own use. They kind of sign a license when they get it that says they're not going to exploit it commercially, but they can use it in any kind of research. We're still debating this a bit for overseas people, but it was felt that this was the bonus that the .ac.uk researchers got from us being funded by EPSRC. So overseas academics didn't get such a good deal. But they could get it two ways: they could either buy the library at a much reduced price, or if they're doing collaborative work with us, we got permission to give them a free copy. They still have to sign a license agreement, but they got a free copy. So we distribute codes to other academics that way. And then it just left the commercial ones, which, frankly we have nothing to do with and just say, "Hyprotech, come on. Please just do the negotiations." Although they do come back to us sometimes.

HAIGH: So in the '70s and the '80s, did you ever have a sense of how widely used it was outside of Harwell? Would there have been dozens of universities?

DUFF: Certainly dozens, and also worldwide—well, worldwide being other countries in Europe and America mainly. Certainly in the dozens. When it got the hundreds, I don't know. It's in the thousands now, but that's sort of cumulative in a sense. So I'm not saying we've got a thousand completely active users. Well, actually, we probably have a thousand completely active users, but not using the whole library, using various routines. We do monitor things a lot more now, because people pick up codes from the web and they have to fill out some data, so that gets recorded.

HAIGH: Prior to this change that you mentioned around 1988-'89, did you have any problems in terms of, say, publishing part of the library in ACM TOMS?

DUFF: Not really. I published both myself and some joint work in ACM TOMS.

HAIGH: Did you have to go through an elaborate process of getting forms, getting permissions, that kind of thing?

DUFF: Not elaborate at all. It wasn't any harder than just publishing a paper. At one stage you had to go through some kind of security. Since we were not working on anything that was remotely secret, whether nuclear power or whatever, that wasn't really a big issue.

HAIGH: So did you have to have security clearance to work at Harwell?

DUFF: I did initially, yes. I was positively vetted, as they call it. There's a term in America, Q cleared, effectively. That was back in the old days. Again, it wasn't from weapons because I wouldn't have done weapons research, but I guess all government establishments at that time were doing that.

[Tape 3, Side A]

HAIGH: Let's return then to a couple of other aspects of the library's evolution at the time. One of those would be adaptation to new platforms. So you said that originally the earliest version had been for the IBM 7090, then you switched to the IBM 360/370 series. Did it stay just based on the large IBM machines or was there work to port it to DEC machines and other kinds of platforms?

DUFF: Yes. I mean, we started to get aware of portability issues I guess sometime in the 1970s. The first main effort was, as mentioned earlier, making it run under PFORT. That was actually a major effort and many codes fell over when we first did that. That effort was happening maybe in the late '70s, certainly after I had come back to the lab from America and Newcastle. So we were getting more interested in portability at that stage. It wasn't necessarily specifically targeting any other machine, and we've not done that too much until the Cray, really, until we had vector machines.

HAIGH: So what was the motivation? Was it because there were more different kinds of machines showing up inside Harwell, or was it to make it easier to share it externally?

DUFF: I guess the main driving force was internal demand, but we were very conscious of the fact that it was difficult to distribute a library that wasn't portable.

HAIGH: So you said that the main technique that you used for that was just to conform to this PFORT portable subset of FORTRAN?

DUFF: That type of thing. That's right. That was a very big tool for quite a while with us. So this must have been in the 1970s before FORTRAN 77 came out, I don't know what year that was, then things would have changed, but up until that point, we were using PFORT on FORTRAN 66.

HAIGH: How about that challenge that you mentioned of the Cray with its vector architecture?

DUFF: Yes. Well, there are two aspects. One is on a personal basis, we should rewind the clock and talk about my visits in the States in the '80s which, were very fundamental in changing my whole orientation on things, and certainly were very fundamental for me in understanding vector computing. So we had Cray machines at Harwell just before we left, notably the Cray 2. And, of course, when we came over to Rutherford, we went from a Cray 1 to a Cray XMP and then a Cray YMP. So we had quite heavy involvement with vector machines fairly early in the development cycle. Well, maybe not that early, though we had one of the first Cray 2s that existed, and we had a Cray 1 at Harwell before that.

HAIGH: Was the Cray 1 at Harwell the main machine that people would use for big problems?

DUFF: For big problems, yes. But I think we still had IBM equipment, quite large machines as well.

HAIGH: So what were people's experiences when the Cray turned up? Were the users very excited?

DUFF: We had quite a lot of hype because they were expensive—especially the Cray 2, very expensive. And it was the last hurrah of centralized computing at Harwell, one must say. We very much highlighted applications which would do well, and there were seminars and publicity both to politicians and to other scientists to say what a wonderful job we were doing, and we did target some things that we knew would work well. The scientists involved were very happy, and generally I think there was quite an interest in it. From our point of view, we spent a fair bit of time trying to make sure our codes would vectorize. Our sparse matrix work changed radically at that stage because we then started looking carefully at what we call frontal and multifrontal methods. This was before the BLAS 3 was developed by myself and others. We recognized, even then, that to make the vector machines work well you needed dense kernels within sparse codes. Our whole structure of algorithms changed because of vector machines. Since then, of course, you'll find that on cache-based machines, which is nearly every machine that you've got nowadays, and on parallel machines, these same techniques may be slightly differently adapted and are just as important. So that was an important redesign philosophy. But that was largely brought about by trying to adapt to vector machines, in practice exclusively the Crays initially. We didn't really have any Japanese machines. We didn't have one of the Fujitsus rebadged by ICL or vice versa.

HAIGH: So then, coming up with new versions of the library routines that would perform well on the Cray, would you say that that was the major challenge for library development during this period?

DUFF: Yes. We're talking about early in the '80s, yes. Sure.

HAIGH: So when did the Cray I turn up in Harwell?

DUFF: I think that must have been around 1980, but I can't recollect. I had certainly been familiar with the Cray I from the visits to the States before, and as I said, particularly working with the one at NCAR that they had at that time.

HAIGH: You had alluded a moment ago to the Cray 2 as the last hurrah of centralized computing. So at what point did the availability of workstations and minicomputers on a more distributed basis around the lab begin to become a serious factor at Harwell?

DUFF: Late 1980s. Around about the time we left Harwell.

HAIGH: Now, talking to people in some of the US labs, they mentioned that with this technological change, it became very hard to persuade people that the centralized numerical analysis group was still needed, particularly with the responsibility of maintaining a library. It sounds like from what you've said that the issues of what your group faced were pretty much independent of this.

DUFF: Yes, somewhat. So without the issues that we faced, which were somewhat more serious for us and for Harwell, I think there would have been some element of that same problem that you mentioned, the US government labs and the decentralization of computing. But we weren't selling just a library, we were selling support and collaboration and consultation. So that would have helped. I'm not saying other people weren't as well, but we were doing this. We probably would have survived, and I think to some extent were surviving in the late 1980s when the decentralization idea was coming through; maybe it hadn't taken full control by that time. And when we came to Rutherford, they wanted our software, but also our research expertise and also our consultancy, if you like, as much as the software.

HAIGH: So were people running the library on minicomputers and other smaller machines?

DUFF: Yes, well, right. To some extent, it would been have run on VAXes even back before decentralization took over. What tends to happen is that a researcher will get to use one or two codes from HSL, and they'll port these with them onto other architectures that they're working on. Within big labs, like our own lab—we may well port the library also onto other computing systems. But once it decentralizes to a ridiculous level, to peoples' workstations, you don't really have a compiled library up on everybody's workstations. So it's the FORTRAN code. We've always distributed the source code. That's another difference between us and NAG, because we can't try and compile it on every machine known to man; we don't have the staff to do that.

HAIGH: So as long as it's in this portable FORTRAN subset, then it is really easy to use.

DUFF: Right. And I should say that nowadays, we're using FORTRAN 90 quite a lot—not exclusively, and I'm perhaps one of the few people that still uses FORTRAN 77 kernels.

HAIGH: So actually, that brings us to another issue in terms of the evolution of the time, which would be the transition between these various FORTRAN standards that you used.

DUFF: Yes. Well, of course, John Reid's been very integrally involved with that, so that's been useful to us to keep abreast of things before they're announced or anything, which has been very helpful. That aspect of John's work is helpful and I think we'd have to say that it was useful, although I was saying earlier I was disappointed he has spent so much time on it. Yes, so we've tended to develop the library as FORTRAN languages developed and we started using FORTRAN 90 when it was available. Not universally, not for all work and for every routine. Now we're using 95 or essentially 2000 and we are tending to use new features, although in a gradual way. Obviously we use array syntax, derived data types, internal workspace arrays, this type of thing. Some generic coding as well.

HAIGH: In terms of releases, I think you said earlier that it was only very recently that you had this idea of version numbers.

DUFF: About a year ago, or just over a year ago, maybe.

HAIGH: So I can see that the last overall release was called HSL 2002.

DUFF: That's a slightly different thing. So the release I mentioned was on every individual routine. For the last year, we got a release number, and they all start at 1.0.0, and as you make changes, you have to change that number. Even the smallest change—even if you just change a comment, you would have to change the last digit of that, anything that changes the physical deck, source code, etc.

HAIGH: But those are specific to individual routines?

DUFF: That's right. Completely so. Now, on top of that we have a release about every two years at present, we're still debating whether this is the way we want it to proceed in the future. So there was a HSL 2002, and there was a HSL 2004. HSL 2002 was the first one to be thread safe. That was the big blockbuster advertising for that, "Get your thread safe library here." 2004 had a number of new routines in it, but it was somewhat less of a change from 2000 to 2002. 2006, well, we have not promised to do it every two years, but in that sort of time frame we might do a new one. In a new release, many of the routines don't change. Then we just add a few and change a few and encapsulate it as a library and make it available. For the .ac.uk users, any code we do goes in immediately, so it's a continuous update in that sense.

HAIGH: Can you clarify your personal involvement with the library over the years? For example, did you receive official responsibility for it when you took over as group leader, or was being in charge of the library a separate kind of responsibility?

DUFF: Well, in a sense, because our group is very small, when I became group leader I was responsible for everything including the library. When we were still at Harwell, the day-to-day management of the library was in Mike Hopper's hands. When we came to Rutherford, I was the sort of *responsable*, as we say in French, but… Well, since John Reid retired, which was quite a long while, six or seven years ago, we made him sort of HSL manager, and he does the day-to-day management of the files. We have only a single point at which material is updated in the master files, which he does. So he's been doing it for seven years as a consultant as well.

HAIGH: So more broadly, you would say it's been a team effort all along?

DUFF: Oh yes, sure, very much so. Even from the word go, even in the days when Mike Powell with his very strong management, one would love to say, even then it was team, yes.

HAIGH: I think those conclude the topics that I'd like to talk about with respect to the Harwell Software Library. Do you have any other points that you would like to make about the library?

DUFF: Well, we got quite involved. I'm just recollecting visits to America, which we could come back to more, but I did have some explicitly to do with software libraries and HSL. Kirby Fong, who used to work at the old NMFE center which then became NERSK and moved from Lawrence-Livermore to Berkeley, was setting up something with Sandia and other labs called SLATECH, or was it even the precursor to that? It might have even been a precursor to that.

HAIGH: SLATECH I think began as a fairly open ended push for greater collaboration in the mid-1970s, and by the late-'70s they would have beginning to deliver a shared library.

DUFF: So we got involved heavily, along with Brian Ford, actually, and very many meetings and discussions. So I had a few trips to the States about that, which were completely software, not numerical analysis research-oriented. I don't honestly recollect for whether we ended up having any of our routines explicitly in the SLATECH library, but we certainly influenced and may have had some, maybe minor routines even, but some routines. We had a lot of routines for special functions that nobody else in the world had for a while, so they were quite popular to have. So what I'm saying is that there was a context in which HSL, and the work for it, was recognized internationally and we were working internationally with other groups on software libraries. Commercially, we actually have money to pay John Reid's consultancy from it and pay the odd visitor for long periods. It's like pin money that we're getting; it's not big money. As I said, the total we get altogether is under ₤100,000 a year every year, which is very useful, but doesn't match our overall budget. It wouldn't pay our salaries.

HAIGH: And in terms of comparing the library today with the other libraries currently available, would you say that the sparse matrix support has remained its strongest area?

DUFF: Yes, sparse matrix and numerical optimization. There is a second library developed by Nick Gould, a kind offshoot of HSL called Galahad. Well, it was called Lancelot before, and then Galahad. Well, Lancelot was a program and Galahad's a library, and he's got many of the numerical optimization codes there. Some are both in HSL and Galahad, and some just in Galahad, the ones solely in the optimization domain. So if you take the two together, then yes, I would say that the strength is in optimization and sparse matrices. Mostly direct methods, to be honest, but we've got a few interesting things from other research in the library now.

HAIGH: So let's conclude this specific discussion on the element of the library and move to discuss another aspect of your professional life during this period, the various visiting appointments that you've held in different places.

DUFF: Sure. Ever since my visit that I mentioned in '71 to IBM Yorktown Heights, I've very much valued collaboration internationally, particularly in the US. When we had a young family, mobility was easy and I could go away for three months or a year. Harwell at that time, as I mentioned, was just wonderful. It was like a university without students, and in fact some people called it such, and they encouraged travel. It wasn't officially a sabbatical, and sometimes they gave me full pay and sometimes I picked up pay from other people. Argonne, for example, paid me when I was there. But they were quite happy, and my bosses at the time (initially Mike Powell, and later John Reid) were quite happy about this experience gained overseas. As I said, my wife wasn't working then; we had young children and we were very mobile. In fact, we had no children on my first trip to Umeå. That was actually very interesting, because I was doing full courses to Swedish people there in the university in sparse matrices. It was the first time I really taught major courses, and to some extent any courses I give now are based on that, though they've been developed a lot since then. I really tried to do the whole of sparse matrix technology, and that fed into the book that was published with John and Al in 1986. That was just a wonderful visit; very friendly people there.

HAIGH: What kind of work were people doing there? Was it a recognized center for this kind of work?

DUFF: No, it wasn't really. I think I was trying to develop the sparse matrix work by giving the course. Axel Ruhe was the most active—he was the head of the department—but he was actually away, and part of the reason they funded me that he was away in America at that time on sabbatical. But the person who was sort of department head in his absent was Per-Ake Wedin, and he was really working in optimization and least squares, geometrical view point. So we did have quite a lot we could talk about, and he was interested in large systems, and some of his work was applicable work there. But we didn't really get a lot of joint research going there. I did get some joint work going more on the graph theory side of things with one of the professors or one of the lecturers there in Umeå, a guy called Torbjörn Wiberg, whose son is now a very good research at Warwick as it turns out. So anyway, that was really just whetting my appetite, in a sense. I was doing more teaching than actual research with the people there, but continuing research with the people elsewhere.

The second visit that I mention here was to a university called UC Boulder. I went there with one child. That was, again, teaching for a term, really. Teaching, this time, graduate students as well as undergraduates, and working a little bit with people like Bobby Schnabel, who is really more in optimization now, I guess, in global optimization. He was one person I met back in the old days at Harwell when he visited Mike Powell at Harwell with his supervisor John Dennis.

HAIGH: So was that why you chose UC Boulder for the visit?

DUFF: No. Actually, in some sense, yes, partly. But partly I toyed with the idea, as I guess many British people have, including yourself, and you've clearly toyed and succumbed, of emigrating. This is almost solely because of the very good activity in my field in America. It's still attractive, of course. I still go there four or five times a year because of that. So my wife and I had thought about emigrating, and we wanted to have a look. We had visited Bolder before—must have been just passing through—and that was one of the areas in America which we identified to be a really interesting area from the point of view of life and living. It also had some good people, not really in the sparse matrix side as such, but in numerical analysis, people we knew. My intention would have been to develop sparse matrix work if I moved anywhere. Several things worked out well, but that was a very unfortunate time in 1980 at Boulder, Colorado. It was rape capital of the USA colleges at that time, and really it was very unfortunate. The place has developed enormously since. It's a fantastic place now, I would say, but it had certain defects. We went there thinking that we might emigrate there and decided not.

We knew California well from my Harkness Fellowship days and from my wife's brother who lives there. That was famous as being, in most people's terms, a "fruit and nutcase" state, and is still a bit wacky and interesting, and was not apparently conducive to long marriage, so we didn't move there. That was our feeling of what happened to people who went to California in that period. New York, I had been to, of course, with Stony Brook and Yorktown Heights, and we knew that area quite well. That was just a bit too miserable in winter, sticky-hot in the summer, and the climate wasn't on, and it wasn't especially a tremendously exciting place, I didn't think. So that's why we went to Boulder, and we'll come on shortly to Stanford again and Argonne. And that's partly why I went to these three places, because they were identified differently as the states we might emigrate to.

I should say on top of that, I made many short visits to Oak Ridge National Laboratory. I'll come to that again in a second when I'm at Argonne. I'll make a mental note of that, because that was when I was offered the jobs.

In Stanford of course was my old friend Gene Golub. Again, I was invited there. I wasn't teaching so much there; I was actually doing some research and work with, not so much with

Gene then, actually, but with other visitors. He's got a very good visitor program, and that's something that's really wonderful, how he does that, and I was working with other visitors there. I did some papers from that stay, so that was quite interesting. We had, by that time, two young children. Hamish, as only just bon then; actually we had a problem and we had to get his American visa before he was born, more or less, which was quite tricky. Nowadays we could have had more trouble, I guess, but then, security was a bit different. I remember we had to send off photographs of him literally as he appeared out the womb to the American Embassy. That was just a terrific visit, really. We thought maybe California was okay, but I think my wife's brother was on to his second wife at that time, so that put us off again.

The visits were useful in terms of developing my ideas for sparse matrices and presenting them to really rather critical audiences—in both cases, all the faculty were coming to my lectures. Although it was undergraduate lectures largely, all the faculty came as well. In the University of Colorado, Boulder, half of my audience were staff. So I got great feedback, and that was very useful for developing the thoughts on it and for helping with the book. The Stanford visit was great from some research aspects.

But the main one, which we went to as a family—my daughter was by this time just going to first grade in school—this was the visit to Argonne. The main hosts there were Jack Dongarra, whom I had known for a long time, and Dan Sorensen (a colleague of his who at that time was less well known, but is now a full professor in Rice, and quite well known), and Jorge Moré (in optimization, so he's not within my field). Paul Messina was the director or head of the department at that time. As I mentioned earlier, this was getting reasonably close to a kind of Harwell environment. We really loved Chicago—wonderful city—and loved Argonne. In spite of all that, we might have gotten a little bit close to, in fact, thinking about getting a job. I don't know if there was one offered immediately at Argonne when I was there.

But when we were visiting Argonne, I was down at Oak Ridge, which I had been for quite a long time, and I was actually given a job interview. I didn't even realize I was being interviewed for a while, but I was, and was offered a job at some incredibly low salary, about half of what I was getting at Argonne. At these days the US lab salaries were completely different. There was no homogenization. I don't know if there is yet, but there is somewhat more of it. And the fact that I was getting offered half of what I was getting at Argonne, they said, "Oh yes, but things are so cheap here in Tennessee that you can buy a whole row of houses for one room in Chicago," which was probably quite correct. But we thought about it and it just didn't feel quite right. There were some good people there at the time. Bob Ward was the manager, and Mike Heath, well, I think he was younger than me, but he was one of the people I was interacting with. David Scott, who then went to Intel, was also there. It was very active. It was a good group. Of course, Oak Ridge has nearly gone off the map of numerical analysis now. But anyway, that was almost the first real concrete "sign this form" job offer in the states.

The Argonne experience was fantastic for one really big reason. It was the first time I met parallel machines. They had a Denelcor HEP. This was Burton Smith's baby, and I got to know Burton quite a lot in the forthcoming years after that. He was a really great guy, a really innovative architect. That was an interesting machine to use. Of course, they had an Alliant FX80 there, and at CSRD, which was David Cook's place down in Illinois. So it was really just a great time to be there. I did quite a lot of joint work with Jack and we wrote up a report on the new computing architectures that were then coming along, and it was actually published in a US

Congressional Government paper, which I'm very proud of. So I've got the seal and stamp of the President, whoever it was, on it, with my name as a co-author on an appendix about parallel computing at that time in the US. I developed some of my algorithms to run on the Alliant FX-80 and to some extent the Denelcor HEP, although we didn't really ever pursue that because it never really developed as a commercial machine. But it was a very good time to be there.

The following year I came for a couple of months, maybe even have been three months, at the invitation of David Cook. Bill Gear was department head for Computer Science at the University of Illinois Urbana-Champaign at that time. I soon discovered that they were head-hunting as well, and I was offered a really nice, attractive job (I don't know whether this should be recorded), but I was being offered over $80,000 a year in 1986 with full tenure immediately—my age would have dictated that a little bit—and some other bonuses as well. So it was a very attractive package. But by that time we had decided that with the two children and the differences in culture and in the schooling, and the fact that we were really quite enjoying this village by that time, we turned it down. But it was a wonderful stay there. I think this about Urbana-Champaign (and I don't want to be rude since this is being recorded), it's a one-horse town—it's got little to it apart from the university. My wife, after two days, said, "I've been around the town three times and I don't know what else to do." I'm used to living near London here, and there's quite a lot in London. The point is, Chicago is too far away from Urbana-Champaign. I was booked twice for speeding trying to get between the two places. In fact I wasn't booked, I was actually let off the hook. But really, that area is very fond in our memories, the whole Chicago area. It's like Glasgow, almost, it's so nice.

We never thought again about emigrating after that point. It would be interesting to know what would have happened if we had taken it. In fact the nearest we came to "emigrating" after that was to Glasgow in 1987.

So I guess we've had many other shorter visits all over and very many places in America, and of course I've gone to conferences as well all over the place. Minneapolis was another place I looked at and was very interesting, but winter is a bit cold at times. That brought me on to the two things, and we can discuss them tomorrow more, but in 1987 at the point of which we could begin to see the writing on the wall at Harwell, I got involved at Strathclyde University as a visiting professor, and I still continue in that role, and also as project leader in this new center that was starting, a European center in Toulouse, France. That, of course, put the death knell completely on the idea of moving to the States; we've just got too many connections here. I don't think I've regretted it, although I love visiting, and I've got tremendous interaction with people in America.

HAIGH: I'll ask you about Toulouse tomorrow, because I think that goes with the parallel computing.

DUFF: Yes, most of my parallel computing has been done in Toulouse.

HAIGH: In terms of Strathclyde, is that a normal kind of arrangement that you're a visiting person, but on a permanent basis?

DUFF: Well, it's renewed. I'm hoping it's being renewed as we speak for just about another three years. So it's only a three-year appointment there. I think different universities handle it in different ways. So it's up to me. Obviously at any stage I can say, "Well, I'm not getting anything from it; it's not very interesting. I want to leave," and they can do the same. So it's reviewed every three years. It's a funny arrangement. I very much enjoy interacting with people and going up there. I give lectures from time to time. I've never given a full course there, although I may well do that in future. We have a very good interaction when I go, but I just go for a couple of days each term. I have some contact with them by email. I'm on their email distribution lists. But they seem happy, and I seem happy, and it gives me an excuse to go back to Glasgow, of course, as well.

HAIGH: So how much time do you spend up there?

DUFF: Well, it's just a couple of days each term, being three terms a year. I have occasionally been up there longer when I've combined it with family visits, but of course I am notin the university every day when I've done that.

HAIGH: Are there specific ongoing projects there that you work on?

DUFF: We're trying to develop some. We've had some sort of projects, but they've not materialized quite the way I would like. The original business in Strathclyde was that they were trying to persuade me to take a chair there, a regular chair in mathematical software. In fact, they designed the chair for me. I didn't realize until afterwards that it was designed especially for me, but it was. And again, being offered huge salaries to go. The main reason I didn't take it was that the department then had some good people (whom I knew then and have got to know better since, one of whom has retired recently and the others are still there) but not in my area. And the department was really deficient in numerical analysis in 1987. Maybe it was my own stupid fault; maybe I should have seen that we could have developed it. David Sloan later became head of department and really changed the whole department and got numerical analysis going, but he was a senior lecturer at the time and was in no position to be dictating things. Although he's a nice guy and I talked to him, I wasn't aware of his capacity for management, let's say. So it really developed. Now Strathclyde's the biggest numerical analysis department in the UK. It's bigger than Oxford, much bigger than Dundee now, which used to be the largest. It's just huge. Bath and Oxford are the two biggest English ones roughly, and Strathclyde's even bigger in terms of established faculty positions. So if this was 1987 and Strathclyde was as it is now, I probably would have taken the chair. It's difficult to tell. But that combined with inertia and the fact that I do enjoy being here, and the family is happy being here—at that time the family was a bit younger. So in 1987, I decided against, and it was at that point they said, "Let's talk about some collaborations and be a visiting professor," and I said, "Oh sure, we'll do that." I've had some collaborations since then.

[Tape 3, Side B]

Beginning of Session Three on the morning of the September 1, 2005. Rutherford Appleton Lab, close to Oxford in the United Kingdom.

HAIGH: Yesterday we covered quite a bit of material. Before we moved on with further topics, I was wondering if there were any other additional points that might have occurred to you with regard to anything that we covered yesterday?

DUFF: I confess, I haven’t really looked over it too much. I don’t have anything at the moment.

HAIGH: So let’s rewind the clock. We’ve talked already about the institutional aspects of your career at Harwell and then later at Rutherford Appleton, about your visiting appointments, and also about the development of the Harwell Subroutine Library. What we haven’t done is talked about your personal academic work and research or about your involvement in the broader professional community after you finished your Ph.D.. So I suppose we need to rewind mentally to 1972 and pick up with your personal research career after you gained your D. Phil..

DUFF: Well, some of that, of course, we have mentioned in terms of visits and interactions with people, which is very important. After my spell at Yorktown and the Harkness fellowship I was quite well then launched on what was essentially a sparse matrix career, sparse matrix technology, direct methods for sparse matrices, which in fact has continued since then. Early on, what we were trying to do was get efficient algorithms just for solving sparse systems because none such really existed before, apart from the few routines already at Harwell. So there was some development there, looking at other algorithms for matrix vectorization other than LU, checking the feasibility (which didn't look good at that time, though there’s been more work done by myself and others since). But really looking around the area of dense linear algebra and seeing what was good to apply to the sparse case. That was my thesis work and continued for a few years beyond.

One contribution which has stood the test of time, and is used more now than before, is some more combinatorial work I did earlier on efficient algorithms for getting forms like a block triangular form or the simple sounding task of putting non-zero entries onto the diagonal, which is really much more important than one might first think. That has really proven to be very useful in later work, both with myself and others. In fact, it led to quite important work with a graduate student at CERFACS on permuting matrices to put large entries onto the diagonal. That work is used almost universally now as a preprocessor for sparse codes.

Early work, as I said, was really on organizing sparse data structures, sometimes complex ones, and developing dense algorithms to use these sparse structures in an efficient way, which is the important thing, trying to avoid n-squared loops (where n is the order of the matrix). Much of that was written up in the book with John Reid and Al Erisman. As I mentioned earlier, one of my earliest papers was a review paper which was in the proceedings of IEEE, and I still get people coming to me to this very day saying that was a very significant paper I had written, and it influenced their life. It was probably because of the very big circulation the proceedings of the IEEE has. It’s probably the most circulated of any of my papers because of where it was published.

So the work on using dense kernels really came about when we were trying to do really efficient kernels originally for vector machines and later on for computers with caches and for parallel architectures as well. The whole theme there really was that the kernels could be handled by dense algebra even though the matrix was sparse. I gave what I considered to be really a quite advanced, forward-looking talk for the first time I was invited to the Dundee Biennial Meeting to give a lecture in 1981 when I talked about the use of frontal techniques. This was pre-dating the dense level two and level three BLAS by some years, but already I recognized that that was important, even in the sparse context. The beginning of the frontal techniques was a simple frontal method which was developed for some theoretical physicists who had picked up some code from finite element people at Swansea. So we developed from there robust and efficient methods for handling sparse matrices using frontal methods. That was in the late 1970s, early 1980. So that was fairly pioneering work, based on work that had been done previously in the engineering community.

But the parallel aspects of that were rather poor, so we moved on and during the early to mid-'80s worked on multifrontal techniques. The initial work was joint with John Reid, and together we developed multifrontal methods. They had been suggested in a technical report that was never published by Speelpenning from Illinois, but his was pretty inaccessible and was not really even written as a scientific paper; it was just a technical report. Really people weren't aware of multifrontal methods in our community until we worked on them. So we worked on that, understood that, and developed it as a practical method for solving sparse systems. And I would say that's probably the biggest contribution that I've made, probably one of the biggest that John made as well.

But since then, I've been working a lot on multifrontal methods, partly with my graduate students, and then colleagues in France, where we now support probably one of the most widely distributed sparse codes called MUMPS, MUltifrontal Massively Parallel Solver. The people in France now take a leading role in the actual coding and distribution of that code, which is now designed for massively parallel machines. But the origins of that were in the rich multifrontal work here at Harwell at the time, and developing parallel versions, which was done purely by myself in about 1986, when I visited the CSRD lab in Illinois the year after being in Argonne. Then I completely re-engineered the multifrontal method to handle parallel systems. That's really the origins of the work, which was related to this development in France. Then the frontal techniques and realization that one had to use dense kernels even before dense kernels existed, is what made me interested in the BLAS development when they came about later on in the '80s.

HAIGH: You had said that in your earlier work in the early to mid-'70s you had really been surveying the breadth of methods use with dense systems and working out which of those could be adapted efficiently for sparse systems. Now, with this shift to these multifrontal methods, were those techniques that would also be used with dense systems, or was this a divergence between the two areas?

DUFF: Well, in a peculiar sense, it would run on a dense system because basically the kernel for these methods is just a dense solver or partial solver, because normally you're only handling part of the matrix. So if you fed the whole matrix then they would work. Now, I'm not saying it's necessarily the most efficient because there's often some data-copying that would be unnecessary if you knew in advance that the matrix was dense. But really the methods naturally would extend into dense systems, because that's the kernel they're using. It doesn't really mean much. For the dense system, you would have a single front all at once, and then you would vectorize it. So it fits in in a natural way. Put it this way: I wasn't redesigning new algorithms for dense systems. What I was doing was taking very careful use of matrix vector multiplication and matrix multiplication as later appeared in the level two and level three BLAS, and I was using these kernels perhaps at the same time or even before people were using the kernels for dense systems. Probably about the same time. And this is all sort of post-LINPACK linear algebra, if you like.

HAIGH: Now, your list of publications is very long. So if somebody wanted to get the best sense, or read the key paper that you published on these multifrontal methods, which one would you recommend that they look up?

DUFF: The original paper is probably still the best, which was [Duff, I. S. and Reid, J. K. (1983). The multifrontal solution of indefinite sparse symmetric linear systems. *ACM Trans. Math. Softw.* **9**, 302-325]. That was the first one. Then we later extended it or ran it on unsymmetric systems, which is also very important because of later developments into the MUMPS code. [Duff, I. S. and Reid, J. K. (1984). The multifrontal solution of unsymmetric sets of linear systems. *SIAM J. Sci. Stat. Comput.* **5**, 633-641].

At the same time, I continued work on general sparse codes and that culminated, or at least more recently culminated, later in 1996 before it appeared as MA48. But before that, I had worked on MA28, and I did MA28 by myself and then I collaborated with John Reid for MA48.

The most recent development in multifrontal (so we're jumping to and from) is a serial code that I wrote, for symmetric indefinite systems that's used in very many optimization packages, particularly, which is called MA57, which is in the HSL. There are several new things in that, and some of them were developed jointly with a graduate student at CERFACS. So this work is still continuing. It has not stopped.

HAIGH: Maybe we should talk a little bit in more detail about the MUMPS package, then.

DUFF: Yes, MUMPS. I had done this initial work on the Alliant FX-80 (and this is back about 1985, '86, right about that time) in Illinois and, as you know, I started at CERFACS—this a center in France. The group I lead is called the Parallel Algorithms Group, which is primarily involved with numerical linear algebra and parallel algorithms, and it's a more applications-oriented center, mainly in climate modeling and computational fluid dynamics.

HAIGH: I'm not sure really if you've told the story of how you came to be involved with that.

DUFF: Well, back in the early to mid-1980s, after the Hong Kong work with Harwell, I became known as a kind of expert on the future directions of supercomputing. I suppose, in a peculiar sense, at that time you could think of me as a European Jack Dongarra. I know Jack very well; I'm sure he wouldn't object to that comment. I wouldn't claim it anymore now, because I've got more interested in the mathematical aspects of various things. So I was really going around the world talking about supercomputers, how they could be used (even if I wasn't using them all, but the ideas). I did a rather large survey for the European Union at the time on supercomputers in Europe, which was half restricted, the publication, because it was commissioned by them, but I was able to produce some versions of it that were more available.

HAIGH: And that's got a 1984 date on it.

DUFF: That sounds right. The work was done, of course, two years before that.

HAIGH: So you were working on that from, say, 1982 onward?

DUFF: That's right, yes. And in fact, in doing that, I actually interviewed every single manager of a supercomputer center in Europe. Now, that would almost be impossible now. First of all, you find supercomputers have somewhat changed, but really there's just so many more centers now. But then there were three or four at a maximum in any country. Germany had the most at that time. So that really was a very comprehensive report for that period. That got me known quite a lot in the European and international community, and I gave several talks to Department of Energy people and others on this work (some of them when I was at Argonne for the year).

HAIGH: In general terms, in this period of the 1980s, what was your impression of how well supercomputing was developed in Europe versus the US or Japan?

DUFF: It was holding its own in some ways, probably more than it is now. The difference now is that you've got these massive machines that are supported by the Department of Energy and its labs. If you took them out of the equation, of course, Europe would probably be almost in a better state than the US, even today. Back in the old days, although the Department of Energy labs and the DARPA projects were supporting high-end computing, there wasn't really such a big investment in single special-use machines. So the companies like Cray, BBN, CDC, ETA, Alliant, etc., were close to having as many sales in Europe as in America, maybe only a little bit less. Certainly they wouldn't have survived the same way just from the American market. Cray for a while was really kept going by European sales. So we were really quite competitive in the sense of the available hardware and in the work being done in the supercomputing arena at that time. And still are, outside of the big iron stuff, the ASCI Blues and Reds, and the machines in the weapons labs, which are much bigger than anything we've got access to in Europe.

HAIGH: Do you think this report had any definable impact on supercomputing policy or future developments in Europe?

DUFF: I think that the European Union then became aware of the fact that, although we running well, there were critical issues in certain parts of Europe. For example, Germany, was really quite well developed at that time, the UK wasn't so very bad, but other countries like France were a bit behind. And so they got a feel for the distribution of supercomputers and what they were being used for, and I think that possibly did feed in a bit.

This is a little bit political and a lot of people might disagree with me, but one of the troubles that we had in Europe was that the European Union put money into trying to support people to manufacture supercomputers in Europe. Now, I'm not saying that these exercises weren't useful in developing expertise and knowledge and techniques. I'm thinking of things like the Suprenum Project in Germany with Ulrich Trottenberg. Honeywell Bull were involved, and they had a separate company that was set up in Lyons with some of my colleagues, mostly from France, which was supported by the European Union to develop a supercomputer. By and large, it wasn't wasted money, but in the sense of developing a supercomputer it was because none of them ever really succeeded. They never even looked like succeeding, and there was a lot of money put in, money that maybe could have been more efficiently used. So in that way, the European Union could certainly have spent money more wisely in developing expertise and using machines, even American or Japanese machines, rather than trying to develop our own. I wasn't advocating for or against this in my report; it was just an informational thing. I think they got overexcited about the aerospace industry, which even then was starting to look very attractive, and in recent years has kept par, or even gone ahead of, the main US aircraft manufacturer. So I think there was some thought that we could do the same for supercomputing as they were doing for the aerospace industry but it never really got off the ground anyway. So the European supercomputing industry works in combination with the American or Japanese producers, and maybe does some add-on things or some value-added stuff to supercomputers, but that's about the limit of the hardware input that's happening today or has happened. So as I said, I don't think it was totally misguided or totally wasted, but it never succeeded in its main goal, let's say. Where were we now?

HAIGH: You brought up that report in the context of becoming "the European Jack Dongarra," and I think that was in response to a question about Toulouse, which was in response to a question about the MUMPS package.

DUFF: Yes, that's right. So what happened was, because I got known in Europe largely from this report, and then giving talks to various people connected with it, I got involved in an organization that was investigating the possibility of establishing a European supercomputing center in Toulouse. This initial investigation was to establish a European Union-funded center in Toulouse, a bit along the lines of some other European-funded centers, one example being Ispra in Northern Italy, which my old friend, George Bishop, who was one of my professors in Glasgow was heading at that time. I interviewed that center; that's when I met him, preparing this report. Unfortunately the people in Toulouse chose a bad time for trying to get this organized as a European center in the sense that the European Union was starting to hemorrhage money and ISPRA was proving very expensive. Maybe some of them might have liked to have close it down, but that was never in the cards. However, the idea of setting up another center with European money never really got off the ground. But the people in Toulouse did get some supporting funds to investigate things. The main person who was driving all this, who is still involved in the CERFACS in Toulouse, was called Pierre-Henri Cros (they pronounce it "cross" in that part of France). He was very energetic. When I first met him, he hardly spoke any English; my French was hardly better, but now his English is certainly perfect.

HAIGH: And your French?

DUFF: My French has not really advanced so much since then. It has a little. In terms of vocabulary a lot, and in terms of my being able to read French documents it has improved a lot. But the center was always conceived of as being a European center and by some sort of definition that meant that the scientific language would be English. On my team, which is the Parallel Algorithms Team at CERFACS, slightly over half are usually non-French, although we have Belgians who are French-speaking, for example. The central language for nearly all written documents is English, and seminars are in English. So it's not really been ideal for me practicing my French, particularly since as one of the few native English speakers (if I can say I am that, coming from Scotland), I feel duty-bound to speak English even more there just so that the French people and the other nationalities (we have up to 17 or 18 different nations involved at CERFACS at anyone one time) will in fact hopefully improve their English by listening to me and talking with me. So my poor old French has got a bit better, I guess, but it's not totally fluent. Speaking it is not so bad, but listening to it I find a problem, particularly with the Toulousian accent.

HAIGH: What proportion of your time have you spent there over the years?

DUFF: Really from the word go, which was in 1987 when I started. I've been working one week a month. So every month I fly down to Toulouse and spend a week. Not all the time, but normally I'll fly during the working days. So the normal pattern is that I'll fly out on Monday, depending on the times of flights, which have varied over the last 18 years, and I'll fly back on the Friday. Sometimes I'll go for less time because I may have an important meeting in England, on the Friday or the Monday, and sometimes I'll stay over until the Saturday or go in on Sunday if there's something important on. But that's the general pattern. The week varies each month, depending on other commitments and what's happening elsewhere and in Toulouse.

HAIGH: How large is the center as a whole, and how large is your team?

DUFF: We have people at all levels. We have sort of permanent staff; we have post-docs with two-year appointments, we have Ph.D. students who are usually three- to four-year stays; we have stagiaires who are people like sandwich students in English who are finishing off their first degree or doing their qualifying for the Ph.D. (it's called a DEA in France) and they'll come and do sort of mini-theses with us. So if you include all these people, and a few administrators who are all scientists (there's some secretarial people, but the other administrators are scientists) then we're up to about 90 now. My team has varied in size from 12, just over ten, up to over 20, depending on these more casual people that are in for two or three years. Right at the moment I think we're running a little bit at the lower end because I've got some unfilled positions, so we're down at about somewhere in the low teens.

HAIGH: And how did the European funding climate develop?

DUFF: Well, there was quite a lot of small money for supporting and investigating the center. That was coordinated through a thing called AECT (I can't remember what it stand for) CERFACS. The name CERFACS was chosen very early on. It stands for *Centre Européen Recherche et de Formation Avancée en Calcul Scientifique*, which means "The European Center for Research and Advanced Training in Scientific Computation." So it's kept that name right from the very word go. We had quite a lot of grants from the European Union, some money from the local French region, some money from the national French government, and we were supported as GIP CERFACS, *Groupement d'intérêt public,* at CERFACS, that was the status there. That continued for a fair while until about 1991, I believe.

Then we were supported largely by governmental organizations. We were a public organization. We had about two-thirds of our money from French sources and then one-third from outside France. We had sort of a funding crisis around then. I think we had been spending more than we had been earning, mostly on equipment. We had a lot of very interesting machines early on, which was actually a great environment and good for publicity. From not long after 1987, before any other Europeans were involved in the SC series ofsupercomputing meetings that were run annually in the US, we had stalls for CERFACS. We got quite well known by people, both because of myself and my contacts in the States and also by attending this meeting and doing a booth.

We had a range of really interesting machines. We had a BBN Butterfly; we had an ETA-10; we had a big Convex; we had an Alliant. We really had a lot of machines that were pretty top-of-the-range, not necessarily in terms of raw computing power but in terms of innovative architecture. We tended to get one of the first in the world, and certainly the first in Europe, of most of these machines. Unfortunately, our budget didn't quite sustain that and we had a bit of a crisis around about '91, and the whole thing was revamped, reorganized. Our status was changed. We're not a *société civile* in French law and we actually have shareholders who are major companies in France. These are really five now. So there's Aerospatial. EADS, the French-European aircraft manufacturer of the Airbus planes, is one of our big sponsors. EDF Electricité de France, is another. The local place where we're actually housed is in Météo-France, so they're one of the sponsors. And the NASA of France, which is called CNES, that's our fourth major sponsor, SNECMA, which is an engine manufacturer, aircraft engines and such like, is also a supporter now. There are various other bodies that we collaborate closely with, although they're not officially sponsors, like CEA, which is the atomic energy authority in France and INRIA which is the big French public body that does research and computing and mathematics. The CNRS, which is another big French research organization, has more recently got involved with CERFACS, mainly through the climate modeling people and other groups than my own.

HAIGH: Did it work as a service center where, if sponsors have got some computing that they need doing, that they can get time on the machines?

DUFF: Well, to some extent, that was happening early on, and a little bit happens now. But in fact, we're no longer on the forefront of the actual hardware side. I mean, we've got some nice parallel machines, and we get some to do experiments with, but we're not running in the same way as the top or most innovative machine in the world; really, that emphasis has changed somewhat. So we actually buy time or get grants for time on public machines. For example, there's a big machine in Montpelier, which is a French academic machine. We have time on there, and we can get time on machines in Grenoble, which are run by CEA, and then the new CEA machine is in Paris which is publicly available as opposed to the one in the weapons lab. We can use these as well. So we now have access to machines in other places.

But we do advise our partners and other people about buying new equipment from our knowledge of it and have, from time to time hosted a machine that they have been looking at on a more temporary basis, to see how it performs before they go on to make a purchase. So that orientation has changed. The applications are stronger now. In the early days of CERFACS—you were asking about the team sizes—my team wasn't quite half, but it was the large, biggest team within CERFACS, and that reflected the fact that we were experimenting with new machines. One of the main things we did there a lot was to try and run the BLAS, the level three BLAS, and we developed a level three BLAS for the new architectures. Some of this work was joint with people in Tennessee with Jack Dongarra's group and others. And Michel Daydé, who was a senior for a while at CERFACS, and then became a professor in the local university and still is there, was very much involved in that work. I should mention that the MUMPS researcher, the main person who was initially my student and then went on and is now a professor in the local university, is called Patrick Amestoy, and a succession of other students that he and I have had have been involved in working on the MUMPS package, so there was a cast not of thousands, but quite a lot.

HAIGH: Let's pull back to the MUMPS package then. Was this something that was distributed and used at other sites?

DUFF: It was under development a lot, and as I said, the first code, the MUPS code, which is MUltifrontal Parallel Solver (not massively, and not intended to be massive, in fact it was for shared memory architectures) came about with the work of Patrick and his Ph.D. thesis that I was supervising. That was initially based on work I had been doing with the Alliant in Illinois in 1985, '86, just before CERFACS started. Patrick was one of my first students and had worked in Houston for a year and knew a little about the NEC machine there at that time. And of course, like all French people, was very well-educated in applied mathematics and mathematics in general. So that was great. His main research was on this MUPS, and together, we worked on that. There was a code that went into HSL from that, a part of MUPS called MA41, which is still there. It's a really good code. It's been developed a bit by Patrick and his wife, in fact, since then. They've made other versions of it.

In fact, Patrick's wife, Chiara Puglisi, was another student from Italy, that came to CERFACS in 1988 and worked on the sparse QR algorithm. She met Patrick at CERFACS and they got married. In fact, we have a habit of being somewhat of a marriage bureau in the sense that several of my students have married each other over the years. Not only in my group, but it has happened in other groups as well. The average age of the center at one time was 27, and that included people like myself who are somewhat older than that and other people who were even older than me at that time. So 27 was quite young, and there was a huge number of people about 25, most of them single, and they didn't always remain so.

So that was MUPS. The MUMPS development was the development fordistributed memory machines. We were involved in a major European project called PARASOL. There were 12 partners in five counties, 2.1 million ECUs (which is a precursor to the Euro, it was an accounting unit at that time rather than a currency). And it was coordinated by the people in Pallas in Germany, who were one of the companies to emerge after the Suprenum Project sort of split up. Karl Solschenbach was the administrative leader of the project, and we had 12 partners of whom two were in management, five were end-users, and five were the code producers and researchers. We had people at MSC Nastran, Polyflow, and various other companies. So it was really quite a mammoth thing, and there was a lot of stresses and strains in this, as you can imagine, and a lot of deadlines which we would have rather not had. One of our targets, then, was to make MUPS massively parallel. So we got a lot of support, and we had several people working on it; I guess the most significant who still works on it is Jean-Yves L'Excellent, who is now in Lyon working for INRIA. He was involved very much in this project. And another key person was Jacko Koster, who has just recently accepted a job to manage supercomputing for Norway, from Trondheim in fact. But he's been working at Parallab in Bergen more or less since he left the RAL-CERFACS collaboration. So CERFACS and Rutherford Appleton Laboratories were two of the academic or research partners in Parasol, which was rather amusing, since I ended up with two hats on all the time.

HAIGH: Is PARASOL like the umbrella?

DUFF: Well, that was the logo, but it was for "PARAllel SOLution". It wasn't just MUMPS. They had multi-grid methods by some of the people in Germany who pioneered multi-grid methods; they had FETI, the main decomposition-type methods of various kinds. So it was a range of different methods for solving large problems, quite often finite element problems, but not uniquely exclusively so, genuine industrial problems in Europe. But also, the symbol for Parasol was a parasol—it was an umbrella.

HAIGH: Were machines like the KSR and Thinking Machines models being widely installed in Europe?

DUFF: "Widely" would be another issue for KSR, I don't think there was so many installed because I don't think that had a very long life before it packed up. But as I said earlier on, CERFACS was one of the most ambitious places for getting the architectures in for a while, and this was in the late 1980s to early '90s—until 1991, really, so from 1987 to '91. But then the reality came that we couldn't actually afford them, so we had to change our track. But there weren't an awful lot of places in Europe getting innovative architectures. There were still people buying big Crays, big Fujitsu machines, big NEC machines, Hitachi.

One thing about not having your own supercomputer industry is that you don't care whom you buy machines from. You get relationships with companies, and Cray has always had a very good relationship with people in high-performance computing in Europe. But you did not buy a Cray because it was American; it was because Cray was a good machine. Some people would buy some of the Japanese machines because they were good machines. We never had any of the very stupid situations that occurred in America. NCAR, when they were forbidden from buying a machine because it was not American. It set back their scientific computing by about five years there. Bill Buzbee, whom I know well, was director there and was stopped for political reasons from doing what was scientifically far and away the right thing to do. So we've never had that really, in Europe. Now and again at a different level we've had it. Like with ICL, when we were meant to buy ICL equipment, even if it was crap, because of the politics. So I've had some aspects of that, but within the supercomputing arena, because we've not really had our own industry properly, just the nibbles of it. But in general, we can buy American, Japanese, or whatever, what we like, buy the best machine that we believe we can get for the money at the time. So that's really what people were buying, rather than the innovative machines that we were just talking about earlier, the kind of BBN-type, Butterfly type things.

HAIGH: I know that in general people found it hard to develop algorithms to exploit these massively parallel machines. Was the MUMPS project as successful as you had hoped it would be?

DUFF: Yes, I think it met its aspirations. It depends on what you mean by "hoped it would be". Obviously, we had feelings for how it would do, and probably it has developed now to what we imagined at the beginning. But maybe at the end of the PARASOL project, which was 1999, it hadn't quite reached exactly what we had hoped, but it was enough to satisfy the European Union that we had done a good job. You know, they hold back a little funding in the end until you do the final review, and that was okay. We got funded. So they were happy enough. So it met the target enough for that. But since then, it has been developed a lot more by Patrick and Jean-Yves as well as by subsequent students that we've had at both CERFACS and at Lyons.

I should add that there's a very close link between CERFACS and a major engineering school in Toulouse called ENSEEIHT (I couldn't tell you exactly what all the letters mean). That's one of the *Ecole Nationales* in France, which are sort of like high-powered, super universities. The *Ecole Polytechnique*, of course, is the most famous, but generally the *école*s are the big, famous universities. And we've had a very close relationship. In fact, nearly all their faculty in their applied mathematics and numerical department have either been graduate students or long-term visitors at CERFACS. So we've had a lot of interaction with them, and with a lot of students through there. And now people like Jean-Yves has had his own students in Lyons who continue to develop this MUMPS package. So it's really quite a major collaborative effort.

We have also had very close recent collaboration, with the Lawrence Berkeley lab, what used to be called the NERSC Lab. The person in charge of the Computational Research Division is Horst Simon. We've had close collaboration with Esmond Ng's Scientific Computing group there, , and in fact, we had a Franco-Berkeley grant for a while that we got with Sherry Li as the main principal investigator in America, and myself and Patrick from the French side. We had a very good collaboration, and then Patrick went on a sabbatical for one year to NERSC, to Lawrence Berkeley. They were developing, and still do, a code called Super LU, which had been part of Sherry Li's thesis with Jim Demmel. So part of what we were doing was looking at the two approaches because they don't use multifrontal, they use supernodal. So different algorithms, but still trying to do this kernel which is dense algebra, so it's still the same idea there, and there was quite a lot of collaborative work. It's still continuing, in fact. As we speak right now, one of my ex-students, whom myself and Patrick Amestoy supervised and finished his thesis about a year ago is presently visiting Lawrence Berkeley now and working with Sherry Li on some other work. So we had very close collaboration, and it was certainly supported by this Franco-Berkeley fund.

[Tape 3, Side A]

HAIGH: So that's PARASOL, MUPS and MUMPS covered. Now, what do you think we should address now? Actually, you still haven't talked about Jim Wilkinson.

DUFF: All right. So I met Jim, as I guess many people did, in the series of meetings which used to be called Gatlinburg Meetings and are now called Householder Meetings after Alston Householder from Tennessee. Basically this is an invitation-only meeting for numerical linear algebraists. It happens once every three years, usually alternating between Europe and the US, or North America, actually, because we've had some meetings in Canada. These are really very intense meetings, very hardworking meetings, but very sociable meetings—a lot of drinking goes on as well, which certainly Jim Wilkinson wasn't shy of doing, nor myself or Pete Stewart or Gene Golub, indeed all the people that you hear about. I got to know Jim a bit there, really. He was a very affable, very bright person, of course, but a very nice guy. Indeed, he was instrumental in some of these promotions that we talked about earlier yesterday, so I needed to get somebody quasi-external who could support it, somebody nice. Of course, since he was the first and for a long while the only Fellow in the Royal Society in numerical analysis, he was an excellent choice for supporting me, and he did that, and helped a lot. Of course, unfortunately, he died in 1986, which was really far too young when he still had a lot to contribute. He had retired formerly from NPL before then because they had to retire at 60. But he did spend almost more time teaching after his retirement, visiting Stanford and giving lectures there, which is something he didn't do in England. As I mentioned, he wasn't really involved in the academic side in the UK. Jim was, of course, a very big influence. His books in particular, his eigenvalue problem book [J. Wilkinson. The Algebraic Eigenvalue Problem. Clarendon Press, Oxford, England, 1965] are major landmarks in our field and a pre-requisite reading for every graduate student in numerical linear algebra.

HAIGH: In general, was there much interaction between NPL and Harwell?

DUFF: Not an enormous amount. There were two optimization schools in Britain in the early days. One was Mike Powell and Roger Fletcher that I mentioned, and the somewhat younger and newer group, which was formed in NPL, was with Walter Murray and Philip Gill. Walter is now in Stanford and Phil Gill is now in San Diego. They both emigrated to the Stanford region quite a few years ago now. There was a rivalry between the two groups. As I said, Mike's a nice fellow, but quite aggressive intellectually. I think he and Walter had swords drawn occasionally. Interestingly enough, Brian Ford has had a closer involvement in optimization through the Walter Murray side, rather than through the Mike Powell one, because we had our Harwell Subroutine Library. So obviously we knew each other and there was some discussion, but there wasn't really a lot of collaborative work. It was more competitive work, almost in a sense. In our community, we're actually very good for being friends and collaborating. The optimization community, which is slightly different from numerical linear algebra, is maybe not quite so friendly. It's getting better I think. But in numerical linear algebra, we did have a lot of contact. We weren't doing any joint research exactly with Jim, but certainly talked a lot. Jim was very interested in sparse systems, but didn't really work in them himself. Obviously he was interested in how any techniques could be developed or extended for sparse systems. The close contact was in these meetings, which are really very important meetings and are still regarded as such today. I still go to them. We just had a meeting in Pennsylvania last May. The next one will be in Germany somewhere near Berlin in three years' time.

HAIGH: So let's talk more generally about your growing involvement with the international mathematical software community.

DUFF: I should just quickly mention that Fox and Wilkinson, of course, were very good friends, and so I did meet Jim through Fox occasionally, and I have been out to dinner with Fox and Wilkinson and people. Mostly after I was a graduate student, because Fox was a little formal, and you wouldn't call him Leslie until after you graduated; that's for sure.

HAIGH: So it was apparent talking about your student days that you were someone who was very much interested in joining groups, being on committees, and starting societies. Did that same kind of impulse carry over into your professional academic career, do you think?

DUFF: Very much so, and maybe even to a greater extent than it should in terms of the amount of time I spent with it! Yes, I got involved with the Institute of Mathematics and its Applications quite a long time ago. I would have to look up my records to see; maybe it appears there. So I was on the council. I made council really quite a long time ago, and I have continued to work with them since then.

HAIGH: So the dates they have here are that you served as the editor of its newsletter in numerical analysis from 1976 onwards.

DUFF: That's still going! That's right. Three issues a year and I'm still doing that.

HAIGH: What niche does the newsletter fill?

DUFF: It's mainly addressing UK and to some extent European numerical analysis, only just because it would swamp me if I tried to cover stuff from the US. So we have a regular list of seminars that are happening in the UK in the next term as it were. Essentially it comes out once a university term. We have a list of technical publications by people in the UK and to some extent European universities over the previous three months or four months. We have a list of visitors that are visiting Britain. We have short articles, occasionally. We have notices about conferences—we have a whole conference section that is international, that is conferences anywhere in the world, and that's quite a large list. We have information on how to contact people. We have regular news from NAG, so every issue NAG will write an article. David Sayers of NAG Central Office supplies me with a copy for that part of the newsletter.

HAIGH: So you've been editing this thing for almost thirty years.

DUFF: Yes, I guess so.

HAIGH: My impression is that with that kind of newsletter, people usually serve about four years and then hand it over. What's the particular satisfaction that you get from it?

DUFF: Well, I don't know. Communicating with people to get the information from them was interesting in itself, and establishing a network that I could get the information from was interesting. I was in contact with people I wouldn't otherwise have been in contact with, because of course this is the whole of numerical analysis and not just numerical linear algebra. So that was interesting. Why did I keep it on to this day? I guess I've really not managed to persuade somebody else to do it, really. So whether it will die with me when I stop doing it, I don't know. I might make more strong efforts to get rid of it. Maybe a thirty-year period would be quite a nice one to stop at. It's not that time consuming, that particular one, because I do very little apart from record a few things until the copy date time, and then it probably takes full day, which is quite a lot, admittedly. In fact, probably less than a full day now.

HAIGH: According to your resume, 1976 was also the year that you joined ACM SIGNUM. Were you ever particularly active in that group?

DUFF: Not especially in that group, no. In fact, it kind of was lumbering along. It's a funny group. I think I was involved maybe a little bit before then, or maybe when I first joined it was quite active, but I don't think it exists anymore. I think it has packed up.

HAIGH: That's true.

DUFF: It was never tremendously active. There was some IFIP working groups in numerical software that were going strong, and I went to a couple of their meetings, but I never got too greatly involved in these.

HAIGH: That's the WG 2.5?

DUFF: Yes, that's right. I never got involved in it really, apart from attending several of the meetings, actually, back in the earlier days of it in the late 1970s, early '80s.

HAIGH: Were you involved in ACM more generally, beyond SIGNUM? It says that you were an "Institutional Designee Member," whatever that is.

DUFF: That was because of getting journals and stuff from them that we had to do that for a while. No, ACM of course was involved in the supercomputing meetings, and for a while I had attended all bar one for the first ten or 12 years. I haven't gone to the last couple for a variety of reasons. But I'm not particularly involved with ACM really, no. I joined SIAM a good number of years ago, but I was really a fairly dormant member for quite a while. I got a little involved in the SIAGs, particularly in linear algebra and supercomputing. I'm still currently secretary of the SIAG on supercomputing.

HAIGH: So how effective would you say those interest groups would be within SIAM?

DUFF: I think very effective, very active. One of the main things that they do almost universally, and they almost have to agree to do it to become an interest group, is that they run meetings. In most cases, they'll run their own meeting once every two or three years, and they'll also run many symposia at SIAM annual meetings and other more general meetings. So they really are in some senses the main technical strength of SIAM in getting people to collaborate research-wise together. And the meetings that they run are among the best meetings in the world in their fields. So the SIAM linear algebra meeting, they had a few in Snowbird in Utah, and they're having one next year in Germany. They're really good meetings to go to, and most of my team here at Rutherford and a few people from CERFACS tend to go to these. The SIAM parallel processing meeting is, again, one of the best of its kind. We had a meeting last year in San Francisco, and we're having one next year in San Francisco. The SIAGs were something that hasn't always been in SIAM, of course, but I think they're really a very good development, and they do focus activity quite well. And most people are members of more than one SIAG. The big new one is Computational Science and Engineering. I can't even remember if I'm a member of that, but I was involved in the program committee for the last CSE meeting, which was in Orlando, Florida. That was a big success. That's a huge group now.

HAIGH: What are your general impressions of SIAM and ACM as organizations?

DUFF: I don't know enough about ACM to say anything too rude, but I certainly have nothing great to say. TOMS drives me mad. They copy-edit it out of existence and change the meaning of things regularly. I mean, I've got complaints at the journal I edit, the *IMA Journal of Numerical Analysis*, about the fact that the copy-editing isn't done as well as they should, but they go over the top with ACM.

HAIGH: That's the publication staff then, not the general editors themselves?

DUFF: I really don't know. I'm almost looking at this as an author rather than an editor. But I could have published quite a lot of papers in TOMS, because we do a lot of software-oriented things because of HSL.

ACM SIGNUM, I was a member for quite a while, and more or less from when you mentioned up until it packed up. But the communications there, the renewal of membership and that, that was unbelievable. For a computing thing, the communications was terrible. That's very rude, because I don't know much about ACM.

But SIAM, I know a lot about. I'm really very impressed. I think it's a well-run organization, and it's got wonderful support from the community. IMA I think is also wonderful, but let's talk about SIAM first. So when I was asked by Tom Manteuffel, who was the president of SIAM, if I wanted to stand for the SIAM Board of Trustees about three years ago I had to think about it, although you see I like to be involved. But anyway, I thought, “Well, that would be a good thing,” and I certainly haven’t regretted it. It’s been very interesting working on the Board of Trustees closely with SIAM. I’m getting to know the SIAM officers, getting to know the SIAM staff, getting to know Jim Crowley somewhat more, and they really are a very good bunch of people. I mean, their hearts are totally in the right place; they are trying to advance applied mathematics. But they have a much stronger numerical analysis bias, at least the people that are involved in SIAM at the higher level. There’s more in numerical analysis than in, say, the British equivalent, The Institute of Mathematics and its Applications. We tend to have more fluid analysis—not even computational fluid analysis, but just straight fluid analysis. So SIAM is a more natural organization than IMA, in the sense that just about every other president of SIAM has been a numerical analyst.

HAIGH: Would you say SIAM has a strong presence in Europe?

DUFF: It’s got a very strong presence outside of the US. About one-third of paid SIAM members are based outside of the US (if you take away the student membership of SIAM which is given away free and you could argue it doesn’t count the same way). That’s really high, and most of them are in Europe—not all, because some are in Australia and Asia. I think there’s about 200 SIAM members in the UK. SIAM has got the special interest groups, and they’ve also got regional groups as well.

HAIGH: Chapters?

DUFF: Yes, that’s right. Regional chapters. One was set up in the UK and Ireland about four or five years ago. I don’t know if you actually physically join that, but I’ve been a sort of member because essentially all the SIAM members who are resident in the UK are members. They have meetings at the beginning of January, which isn’t an ideal time for us Scotsmen who are still celebrating the New Year, but I’ve been to a couple. They don’t do too much else than that; they sometimes sponsor a talk somewhere or this kind of thing. That’s really just to try and help SIAM members get more out of their membership in the UK, and that’s working well. But it’s not an organization like the IMA or the like the LMS, which I have joined more recently, which is the London Maths Society, originally the pure maths.

HAIGH: Did you have much personal contact with Ed Block?

DUFF: Yes. I knew him quite well. As I said earlier, all this is being monitored and that. He was really a fantastic thing for SIAM and really got it going. The investments are still doing very well, and he got them all organized and set up. There could be a debate about whether he stayed longer than he should, but this is a very natural thing, and it was same with IMA as well. But SIAM owes him a very big debt of gratitude. He was a very strong-minded person. I was never was involved in SIAM at a level to cross him or argue with him. I don’t know what would have happened there. All I know is that Jim Crowley’s a very easy person to get on with, and very efficient, and I really enjoy getting on with him. I suspect I might have had more arguments with Ed if I had been involved at that level at that time.

HAIGH: You mentioned that you are now one of the SIAM trustees.

DUFF: That's right. I just mentioned earlier in the interview that Tom Manteuffel had asked me to stand, and I did stand, and I was elected as a trustee. I'm in my third year now as a trustee, and I'm standing for a re-election, which is happening this autumn.

HAIGH: So what kind of responsibilities do the trustees have? Is a hands-on kind role?

DUFF: They're the board that ultimately runs SIAM in a sense, and they certainly do make the financial decisions. The council of SIAM makes the academic decisions, the policy decisions on things, but it always has to be at least rubber stamped (and it might even not be rubber stamped at all if it was something controversial) by the board and the board financially manages SIAM. So it's the ultimate authority, and it is a bit smaller than the council.

HAIGH: Let's return to the Institute for Mathematics and its Applications. American readers of the transcript might not be well aware of this group, so could you describe it?

DUFF: Well, of course it was preempted a bit by a much newer organization in Minneapolis, which called itself the IMA in America, and that does cause a lot of confusion. But the Institute has been going for forty years. We just had an anniversary. Just like SIAM to some extent, it was formed because at that time the feeling was that the mathematics organizations, in Britain the London Maths Society and in America the American Math Society, weren't supporting applied mathematics properly, and in fact, applied mathematicians were regarded as second-class citizens and intellectually inferior et cetera, et cetera. And so in the case of the IMA, somebody called James Lighthill, who was a fluid dynamicist of world renown—a genius you could say, really, one of the top guys of the field—was sufficiently unenamored, I guess, by LMS (I don't know the fine details, there were a few things written about it in our bulletin) to form, with some like-minded people, this Institute of Mathematics and its Applications, so that the applied maths community could have a stronger voice and be regarded in a better light by their own peeers. It took a year or so to get it off the ground. So it was definitely a split off because LMS wasn't functioning in their favor. SIAM, as I said, has something rather similar for why it was formed, and at not too different a time, I think. The very interesting thing now in Britain, is that the membership of IMA for applied maths is actually over twice as big as the LMS membership, which is I think one of the few countries in the world that this is the case. AMS is much bigger than SIAM, and in most cases the pure maths organization is larger. LMS is very much a Learned Society and very much based in the university environment. IMA is both a Learned Society and a professional organization which has in its members many industrial people and many school teachers of mathematics. So it has a much broader spectrum.

What's happening right now as we speak is a thing called The Framework Study Institute, or FSI, which is looking a ways in which LMS and IMA can collaborate and cooperate more closely. This could even lead to a merger of the two organizations. It's not got as far as that yet, and in fact, it's not the primary thing on the agenda, but it's one of the possibilities that's being discussed. So this is interesting. I think it's the recognition that the pure mathematicians realized that applied maths isn't just people doing crossword puzzles, and the applied mathematicians realize that pure mathematics and applied mathematics are all mathematics really. And the government also wants a single voice of mathematics in Britain. That's caused the Council of the Mathematical Societies, CMS, to be formed, which I'm a bit involved in, because I'm the Vice President for Learned Society Affairs in IMA. I have been for a year and a half. My term goes until December of this year. In that capacity, I have been involved in some of these collaborative bodies that are trying to talk to the ministerial level in the government, and also with EPSRC, the main funding agency, the kind of equivalent of NSF for the UK. So we talk at the highest levels to them as well through this Council of Mathematical Societies, which is the combined one. It's partly this contact that's made us realize that the LMS and the IMA have got an awful lot to contribute to each other, so that's what set up this initiative.

So that was a small aside. IMA had a person called Norman Clark who was the secretary and registrar when it was formed, a bit like Ed Block. I think Norman did have a background in mathematics but he was administering it all. So we have a president as well who is a mathematician. Presently that is Tim Pedley from Cambridge University; he's a fluid dynamicist. Like Ed, Norman was a strong character who really guided it. In fact, I think in some cases there was a question of whether the president was really having as much say as he should. Norman, a bit like Ed, kind of maybe overstayed his sell-by date a little and was eventually more or less sort of squeezed out. It was a bit awkward. I was around about that time on my first stay on the council. Then Catherine Richards took over from Norman for a couple of years, and was more innovative in many ways, but she unfortunately had cancer and died in office. We had two other registrars since then, Adrian Lepper did a wonderful job in sorting out the finances of IMA, which weren't in very good shape at that time. He's still around as a consultant. Very recently, about a couple of years ago, David Youdan became our secretary of the IMA. Now that the finances are fine, he's looking at ways in which we can spend money and use money for the benefit of the community. It is really quite exciting now that that's all been sorted out.

HAIGH: Is that the same position that would be called executive director in an American society?

DUFF: Yes, that's right. I'm trying to remember his title exactly in England, but it's effectively the same thing.

HAIGH: In terms of your own involvement, I see that you were made a fellow in 1984.

DUFF: Yes. So it's a bit different from SIAM. We have grades of membership in IMA, and there's a membership committee that reviews the applications for the different grades. To be a fellow you have to have so many years post-doctoral experience at a certain level after being an active researcher or teacher perhaps, involved in mathematics. They look at your CV essentially, and judge whether you can be made a fellow. It's not a handful of people that are fellows. I guess there's between one and two thousand fellows out of a membership of 6,000 roughly.

HAIGH: So it's not like being an ACM or IEEE fellow?

DUFF: Not that type of thing, really, no. I'm a life fellow as well, but that was because I opted to pay a life membership some years ago.

HAIGH: Then from 1985 to 1988 you were a member of the Council.

DUFF: Right, a member of Council. I've been on and off member of Council for several times. That was a period I was there.

HAIGH: And in 1996 to 2001 you were also Chairman of the Program Committee and Honorary Secretary.

DUFF: That's right. Five years, that of course was for the honorary secretaryship.

HAIGH: What does an honorary secretary do?

DUFF: Basically you countersign the accounts, and you're involved on the Executive Committee and the Council. The Executive Committee is probably a little like the trustees in SIAM. So you're involved in these and the decisions at that level. At that point the chairman of the program committee was sort of almost ex officio secretary; the two sort of came together. So that was a five-year thing. Conferences have always been an awkward situation both with SIAM and IMA, and for the same reasons, and that is that running conferences is not an easy task to do efficiently and cheaply. If you run a little conference locally in your own university, you can make it very cheap sometimes because your room is free and various things. As soon as you start organizing a bigger conference with multiple organizations in a third-party place, the costs start to mount and the registration fees go up. SIAM members have been complaining bitterly about SIAM fees. I would say they have been doing it more recently than in the past. In the past we had trouble with IMA. So being in charge, or the Chairman of the Program Committee, was not actually an easy task. You had to justify and argue and try and get cheaper prices. This was against the background of an IMA, which was trying to balance its books at the time with Adrian Lepper, so I certainly wasn't getting many goodies sent in my direction. That has improved a lot. Now we've got a balanced budget, and I would say that we are at least as good as SIAM, if not even better, in how we can keep down prices at conferences. It's certainly not too dissimilar. IMA tends to run conferences in universities rather than in big hotels, so some of the attributed costs are less, although we have to pay for audio/visual and all that kind of stuff.

So that was what it was for five years, and I got quite involved with the IMA. We had a big strategy weekend meeting that was rather important, and I was involved in that. That was to try and define our way forward. It caused changes in the management structure of the IMA—much more streamlined executive, and things. It increased, if anything, the power of the president as opposed to the registrar secretary, the CEO. I had a few things, to say about conferences that I guess helped.

Then more recently, about three years ago, I became Vice President for Learned Societies in the IMA, and therefore I was also ex officio essentially on Council and Executive Board. Which is why I've got a lot of connection with LMS, because LMS a learned society. So in that aspect there's a good overlap of some of the activities of LMS with some of these activities of IMA. So because of that vice presidency, I've been quite connected with some discussions there. That comes to an end in December, but I've just been re-elected onto the Council. I'm just a Council member as of January 1 next year for IMA.

HAIGH: So have you had any other significant involvement with mathematical software groups or professional societies that we should cover?

DUFF: Well, yes. There's a group that's not really a society, but a society of societies called ICIAM. It used to be called CICIAM, the Committee for International Congress in Industrial and Applied Mathematics and now it's the International Committee for Industrial and Applied

Mathematics. This body consists of representatives of most of the major applied maths organizations including both IMA and SIAM. I've been on the board of that as an IMA representative—not a SIAM one, but an IMA one—since back in the early 1990s. I was certainly a member in '95, maybe before then. So I've been quite involved and I am still involved heavily with that organization, which has certainly been going from strength to strength. We've got more funding than we used to have, and we can undertake some initiatives, other than just the quadrennial congress. The last one was held in Sydney. The next one will be in Zurich, and the one after that will be in Vancouver.

HAIGH: Nice spots!

DUFF: That's right. People bid for the meetings partly, so the competition last time was between Vancouver and Beijing. That was, of course, a very good competition. Both were giving very good offers, but we chose Vancouver. So I was quite involved with that. I am also a member of SMAI, which is the French organization because of my involvement in CERFACS.

HAIGH: And that's S_M_A_I?

DUFF: Yes. That's the French. And they're also involved ICIAM. But I'm not involved in the councils or anything of SMAI. I've got quite enough to do without that. I also joined SBMAC, which is this Brazilian society, because I visited Brazil a couple of times, and Venezuela a couple of times, and Mexico a few times. I had a spell in the 1980s when I was going every year to Central or South America. I got quite involved with people there, so I joined their local organization, which is called SBMAC. I still maintain my membership, but the communication isn't so great, as you might find when you're working with Brazil.

HAIGH: So that's your list. Is there anything else on there you want to mention?

DUFF: I'm just looking quickly, bouncing through this. Some of them were very short-lived things really and not worth pursuing. I was involved at various levels in the UK government bodies that were looking at advanced computing. This was back when I was at Harwell. I haven't done so much since I have been here because there are other people here that have been heavily involved, and I believe I'm not looking for these tasks directly. I had a spell on the scientific council of an organization in Normandy, northern France.

Yes, there is something I do want to mention here, that's the Fox Prizes. I mean, they were really interesting things. I was involved right from the word go in these. The idea of the Fox Prize really came from Gene Golub. This was when Fox was still alive. See, Jim was regarded, justifiably maybe, as one of the most famous numerical analysts in the UK, and was a fellow of the Royal Society. But in terms of influence in numerical analysis and in producing the next generation of numerical analysts, Fox was far more influential because he was the head for many years of the major lab in numerical analysis in the UK at Oxford, and he either directly, or indirectly through his other staff, was responsible for about half of the numerical analysis output of Ph.D.s in Britain for a fair while. I may be exaggerating slightly at half, but certainly a very large percentage. So Gene, who was very friendly with Leslie—most people were very friendly with Leslie—decided that we should institute a prize, but of course it wasn't really up to him to do it. I think he made a nice big donation to the fund, but otherwise the organizational side was left a bit to us. Being a close friend of Gene's, he mentioned it to me and to Nancy Nichols at

Reading, and had also been a student of Lesley's a little earlier than me. We got together and we decided that in order to get this thing off the ground, we should have to get Mike Powell on board, as he was the other FRS from numerical analysis at that time. We collaborated and contacted him and drew up together guidelines for this prize. We've set up a committee. We chose senior people at the time. I wouldn't regard myself senior at that time; I was still quite young. We got that off the ground and it's continued ever since. There's been some very significant winners of the Fox Prize. I think the first winner was a young fellow called Nick Trefethen, who is now a professor in Oxford, and in fact is the newest Fellow of the Royal Society in numerical analysis. That happened this year. Nick Gould, who works with me here at Rutherford, was also a first prize winner. Jennifer Scott was a secondary prize winner. So there have been a lot of people that have come through and have had significant careers since then. It's a very well recognized prize. Very underfunded, and it's still underfunded really. It's organized through the auspices of the *IMA Journal of Numerical Analysis*, which is the journal that I'm the chief editor for and have been for the last 15 years or something, I guess.

HAIGH: So your personal involvement with that has been serving on or chairing the adjudicating committee?

DUFF: Yes. I didn't do that in the first instance because, as I said, I wasn't really a senior person then. I was too old, I think, for actually being an applicant for the prize. You have to be under the age of 31 on December the 31st of the prize year. But later on, yes, I was involved. It's a three-person committee. And the last year you're the chairman. So you take off one person each year and bring on a new one; that's how it works. So that's still going strong. We'll be discussing the appointment of the next adjudicating committee member this week at our IMA numerical analysis editorial board. The chairman will normally be the most senior left of the two from the previous one. So that's still going very well and is a feature of British numerical analysis. Well, world numerical analysis, in fact, most of the winners recently have not been British people, which has been causing some anguish because we haven't had the applications coming in from Britain. It's not that they have been bad or not as good, but we just haven't had the numbers coming in. In fact, some of them have been non-British, but have done their Ph.D. in Britain, so that counts.

HAIGH: Do you think this would be a good time to stop for lunch then?

DUFF: Yes, it may be a good time, I guess.

HAIGH: So this is the conclusion of Session Three.

[Tape 4, Side B]

Beginning of session four, afternoon of Thursday, September 1, 2005.

HAIGH: So I think that we concluded the discussion on your involvement in these various professional societies and groups. But there are several topics remaining. One of those that I know you wanted to talk about was your involvement with test matrices.

DUFF: Oh yes, right. So from the early days at Harwell, around about the time I arrived there—in fact, when I was a Ph.D. student there—it was recognized that you have to have something to test sparse matrix algorithms against. There were very few test matrices—well, none around at that time, really. Fortunately we got a few matrices from Alan Curtis from his ordinary differential equation work, and other matrices came from Yorktown Heights from the first meeting that they had there in 1968. So what I did originally was to collect matrices together into a simple data set and use them to test our codes. But other people who were working on sparse matrices needed test examples as well. The initial set was very small and didn't get distributed, but we built from this the Harwell Sparse Test Matrices. John and I distributed this quite widely. Then later, with our colleagues from Boeing (I haven't really mentioned work with them, but this was John Lewis and people in Al Erisman's group) we got together and developed a Harwell-Boeing Sparse Matrix Collection. This got a little bogged down in the details of just how we stored them, how we presented them, and how we had the data organized in directories and this type of thing, but that's now the main data set that we distribute. Just recently we've been involved with CERFACS and with the people in France in a web-based structure for handling sparse-matrix queries, for making sparse codes available over the web, and within that also making available sparse matrix data sets. So indeed, within that as well we're developing access to these data sets. Since then, one or two other people have developed data sets. Tim Davis, who was an ex post-doctoral student of mine at CERFACS and now in Florida, works with a database of sparse matrices, and the people at NIST established the Matrix Market. So there have been a few other initiatives since then, but for many years, Harwell was actually more famous within that community for test matrices than for the subroutine library. So we did a lot of distribution. Of course, there was no sense of trying to make money or sell them, so that avoided any commercial problems.

HAIGH: I saw references to the "Harwell-Boeing standard" for representing the matrices.

DUFF: That's right. That was part of the exercise indeed. So we had a paper in TOMS that was devoted to that. This also described the Harwell test set [Duff, I. S., Grimes, R. G., and Lewis, J. G. (1989). Sparse matrix test problems. ACM Trans. Math. Softw. 15, 1-14].

HAIGH: So did that representation become a standard?

DUFF: Yes. It wasn't rocket science. It was roughly the kind of data structures that are used commonly for sparse matrices. Matrix Market stores them in two ways. It's called the Harwell format or Harwell-Boeing format. So there's header information, and that's the new part, and then the rest of the matrix is stored in a similar way to how it's stored for use in codes, which is to store only the non-zeros and store them organized by rows or by columns with indices that say where they appear in the overall matrix. So that structure is one of the most common ones, which is one that we experimented with back in the early '70s. There are various structures that we experimented with in the '70s, and that turned out to be one of the most efficient for many purposes. So that was fairly standard in a sense, fairly economical and fairly easy for accessing and printing the matrix. The header information was for identifying the matrix and was designed originally for FORTRAN use. So there was quite a lot of organizational effort that went into it, and a few original thoughts. Not exactly high-powered research, but I guess quoted a lot in some of these papers.

HAIGH: In terms of distribution of the matrices, I imagine that at some point in this process the Internet would become very important.

DUFF: Yes, of course. I mean, that's quite amusing. If you go back to when I originally started distributing matrices, it was things like very big tape reels we would send out. Then we used to send out floppy disks of various kinds, but they were a bit smaller and of course couldn't hold much in the way of sparse matrices. Some of the data sets were very large. I mean, we're talking here in gigabyte terms more than anything else. Distributing through the Internet is all very well, but you have to have pretty high bandwidth lines to handle gigabytes of data. Well, you have to put in quite a few CDs as well if there are gigabytes, but yes. But that's part of this development in France, what we call the TLSE project (not for Toulouse, as you might think, but for Tests for Large Systems of Equations) a system that we've got on the web. There has been a lot of computer science software effort put in that I find very obscure. Not because it's computer science only, but because it's all in very detailed technical French, which is not the technical French that I know from mathematics, so at the meetings I'm not exactly able to contribute as much as I would like to that. But the core of it is to have sparse solvers involved, and one of the other things is to have sparse data sets, so I've been able to contribute there.

HAIGH: So is there anything else you want to say about the test matrices?

DUFF: Not really. It's just that people still refer to the original sets, even when they're using matrices that are not from the original sets, but there we are.

HAIGH: We should also talk about the book that you wrote with Reid and Erisman. [Duff, I. S., Erisman, A. M., and Reid, J. K. (1986). *Direct Methods for Sparse Matrices.* Oxford University Press, London].

You've mentioned it in passing a number of times, but it would be interesting to hear the story of how that came about and how you feel about the results.

DUFF: Originally, it was John and Al that were working on it. They had met at a meeting in Cambridge, must have been back when I was a graduate student, about 1972. And sometime later in the 1970s, it must have been about the time I had the survey paper in Proceedings of the IEEE, I was being asked by various publishers to write a sparse matrix book, which sounded attractive, and it became clear to John Reid, who by that time I was working with in 1975, and myself that that was a bit daft and that it would be much better off to pool our expertise and to be involved in the book with him and Al, which hadn't progressed very much by that stage. It didn't immediately progress very much, even with me involved, but eventually by 1986 we had it published, so it was ready about a year before that, I guess. And that was with OUP.

It was a very interesting exercise, and at that time, I'm sure that almost anything that was known about direct sparse matrices was included. It really was very comprehensive. There was an enormous amount of work done then, and there has been so much more done since. But it was rather comprehensive. It came out about the same time as the book by Alan George and Joseph Liu [Alan George & Joseph W.H. Liu, Computer Solution of Large Sparse Positive Definite Systems, 1981] which was a Prentice-Hall book. But that was just mainly graph theory, and just for symmetric systems, so there wasn't a lot of overlap, and they were fairly different books appealing to slightly different audiences. We were trying to appeal to mathematicians and engineers and users, and they were more in computer science, I guess. So yes, it was a big success. It was a best seller in, some sense, for a book of this kind. We had a reprint quite quickly. Then we had a sort of quasi mini-second edition—it wasn't really a second edition;

there were a few changes we made some years later. And they also brought it out into paperback, because originally it was hardback.

HAIGH: Why did it take so long to finish?

DUFF: It's just that there are other things to do. You should ask about the second edition; that might be a good example. So the second edition, I mean this is 1986, so we're talking 20 years ago. An enormous amount has happened since then. A lot of it, what we talked about in 1986 is completely valid still, so we thought, "Well, all right we'll graft in a bit, probably increase the size. We'll try and drop out a few things that haven't progressed the way we thought they might and have a second edition." We're more than half way through it, no question, but there are still a few crucial things. One of the big issues in writing, the thing that's causing the problem is the section on parallel things. This was about the time of the Alliant FX-80, that was in 1985, so usable parallel machines were only just appearing, and we've got virtually nothing on parallelism on the original book. So there's an awful lot since then, both stuff that I've been involved in and done myself and stuff from other people. So I was hoping this last month to be writing more of that chapter, a new chapter in that, but it's really a very difficult decision as to what to put in and leave out, as always the case in books.

Once we get that chapter done, I think we're into the last 20% of the first draft. But the reason it takes a long time is because new research is more exciting. So if you add all these horrible numbers of administrative things that I've got into with IMA and SIAM and all these things to administrative stuff of being head of the group here at Rutherford and being Project Chief in Toulouse, and all the travel I do for conferences, sometimes it's a wonder that I get any work done! The issue is that the the research that I do takes precedence, of course, over doing the book, so then there's no time left for the book, basically. So that's the trouble, and it is a pain. I think John's a bit nuttier than me. We almost have disputes over it as to who's the worst. I mean, he doesn't have anything like the same range of activities because he's retired now, but nevertheless, he would almost prefer to do anything else other than write a book. And in a little way, it was the same in the first version, even if I personally had rather fewer other duties at that time. But it was a good experience, writing the book, and I think it was a well-recognized and good book. It's still used, even 20 years later. But of course there are many areas that need improving or updating.

HAIGH: Do you have any stories about response from readers?

DUFF: People use it in courses in the US, people like Beresford Parlett. He's used it and been very happy with it, and he's certainly a very discerning gentleman and would have certainly let us know if it wasn't good. I've really not heard too many bad remarks, other than the fact that it's a little out of date. We did approach a number of researchers at the time with drafts or sections that involved them, so we got some feedback from people. I imagine that we'll do that a lot this time once we eventually get this first draft through, because naturally, people look immediately for reference to themselves, and if it's not there they'll tell you, and that may in fact mean some changes.

HAIGH: And you're also involved in editing a tribute volume for Wilkinson?

DUFF: Yes, that's right. This was very soon after his death. *Linear Algebra and Its Applications* wanted to do a commemoration of his life and work and myself and François Chatelain and Jack Dongarra were asked by Hans Schneider and Richard Brualdi to edit this volume. [Chatelin, F., Dongarra, J., and Duff, I. (Special Editors). (1987). Special Volume in Memory of James H. Wilkinson Linear Alg. and its Applics. 88/89]. We were given a completely free range and we were left to do it free hand, and it was quite a big job, because you can imagine that a good number of people were keen to submit papers in honor of Jim. All of the papers were good work really, but some of them were very detailed and there was quite a lot of both refereeing and editing, which we had to handle. We would send it out to referees, but we had to handle all that. So that was quite a big labor, but I think it was a good job that we did. It is ironic that the papers were originally submitted when Jim was still alive and that he died before publication. I also had two papers in a memorial conference that they did for Jim at NPL. That was, again, quite a good meeting and I was sure that Jim would have been happy to have attended if he hadn't been dead by that time.

HAIGH: Should we talk about your work on various versions of BLAS then?

DUFF: Yes, well that's another thing that one gets almost sucked into. As I mentioned earlier, before the BLAS I got involved with kernels using dense linear algebra within sparse linear algebra for efficiency, which were effectively BLAS kernels, or became such. We had several meetings, mainly with Dick Hansen rather than Chuck Lawson involved. This was at Purdue and Illinois. I think we had some of them started when I still around Argonne during the early 1980s. There was some discussion then about what should be done to augment the original Hansen and Lawson BLAS. Now, if you remember the LINPACK package, available in 1979, was developed partly around BLAS 1. But although they do vector arithmetic in BLAS 1, it doesn't work on vector machines, really, because it's too small a granularity. It's what we now call level one BLAS. So it was recognized that the performance of LINPACK wasn't great on vector machines, particularly, and on parallel machines it didn't get anywhere, of course.

So it was recognized that higher-level kernels would be involved. That would be matrix-vector multiplication or matrix-matrix multiplication. And the vendors were quite up for it because they knew how to do it, roughly. These were some of the easier operations because they're such regular operations to do efficiently on then-modern architectures. Mostly they just had vector, with parallel coming in during the 1980s. So we worked hard in various committees to do this. I'm not quite sure how the authorship was fully determined, because in some senses I was quite involved with level two BLAS, maybe less with level three because I missed a couple of meetings. But anyway, Jack Dongarra was probably the main driving force, to be honest. A group which included Dick Hansen and Jack and Sven Hammarling and Jeremy Du Croz formulated the level 2 BLAS. I was in discussions, but I didn't actually agree wholly because I thought there were far too many subroutines and I thought it was a bit complicated. That wasn't the reason why I wasn't involved in authoring it, I'm sure. They had decided in advance. So they went ahead and did level two BLAS, and almost at the same time, we were also discussing in these committees the level three BLAS, with which I was more excited because I was using them already in my sparse codes. We then worked on that with Sven, Jeremy, Jack, and myself.

HAIGH: So how were these committees organized? Were people just invited to take part?

DUFF: Yes, I think it was essentially by invitation, but it wasn't exclusive. I think they were sort of publicized and you could more or less invite yourself, if you wanted. Anybody who was willing to help and communicate and contribute by writing bits of code, or even by discussions and designing algorithms, was very welcome. Jack Dongarra was particularly good at this sort of thing, and he's done it since then for involvement with MPI (even more with PDM, I guess) in various major software projects that he's coordinated. He's very good at trying to get involvement of people at a full level, and trying to have it non-exclusive as well. So it was by invitation for people that you knew that were working in the area, but it wasn't exclusive. I was not fully involved with what was happening on the BLAS, and suddenly Jack suggested it to me "Right, let's do a paper on it," so that's when I knew I was included. The detailed coding was very much done by the people at NAG. Sven Hammarling and Jeremy Du Croz. We all did some, but they were very much involved in the testing as well.

HAIGH: Are you talking there about level two and level three?

DUFF: No. I guess they were very heavily involved in level two as well, but I'm talking about level three, which is what I was involved in myself.

HAIGH: How closely integrated was the level three BLAS with the specific needs of LAPACK?

DUFF: Well, in a sense, LAPACK was being discussed by the same people. It came up in our discussions before it arrived, roughly at the same time. So the two were sort of integrated in a rather reasonable way. Because it was recognized that high-level BLAS were needed for more efficiency, it was recognized that something more efficient than LINPACK was needed, and that's why LAPACK came about, and then it was designed by the level two and level three BLAS. That's how it all happened with the original LAPACK and there have been several revisions since then. I wasn't too much involved with that. That was really getting into a level and depth of programming that I wasn't so interested in because I wanted to spend any such effort in programming on the sparse side. So I didn't get involved in LAPACK, but I did do some programming for the BLAS and some testing, of course. Since then, as I mentioned with Michel Daydé in CERFACS, we were helping with international efforts for testing the BLAS on many different processors and machines. As I said, we had very interesting architectures at CERFACS: testing them, comparing them, tuning them, modifying them. So there's quite a big effort on that over some time.

HAIGH: When you mentioned that you were an author of the paper, is that this one here [Dongarra, J. J., Du Croz, J., Duff, I. S., and Hammarling, S. (1987). A proposal for a set of Level 3, Basic Linear Algebra Subprograms. In Parallel Processing for Scientific Computing. Edited by G. Rodrigue. SIAM, Philadelphia, 40-44].

DUFF: That was a proposal. That was our initial proposal after some meetings that we formulated. So that wasn't actually published. That was the preliminary one, which we put out to a wider community so we could to get feedback, which is another characteristic of Jack's projects, that he's very good at distributing drafts for comment and then reacting on the comment, so we did that.

The two papers that appeared in ACM TOMS were where we had finalized the design and everything. One is a sort of paper discussing the design of the algorithms, and the other is actually the software itself. So they're a pair; they come together. [Dongarra, J. J., Du Croz, J., Duff, I. S., and Hammarling, S. (1990). A set of Level 3 Basic Linear Algebra Subprograms. *ACM Trans. Math. Softw.* **16**, 1-17. Dongarra, J. J., Du Croz, J., Duff, I. S., and Hammarling, S. (1990)] and [Algorithm 679: A set of Level 3 Basic Linear Algebra Subprograms: model implementation and test programs. *ACM Trans.Math. Softw.* **16**, 18-28].

HAIGH: Software, was that kind of generic FORTRAN implementation?

DUFF: Yes, it was FORTRAN. Portable FORTRAN, let's put it that way, rather than generic.

HAIGH: So how would you say that the proposal there in 1987 was received? Did the community embrace it as it was, or were there significant changes?

DUFF: There weren't significant changes, but there was one or two comments made. I was one of the ones arguing, I don't know if I was the only one of the four of us, that I wanted the BLAS 3 to be slimmed down to really encourage people to spend the effort in tuning them for different architectures, either the vendors or people using the machines; in contrast to the BLAS-2, where they had a very great number of routines. If we had too many routines, I argued that that would be counter to people wanting to do good implementations. So we had many fewer routines in level three BLAS. But the issue then was that some people commented that they would like other features. We didn't take them on board. I don't know how many were mentioned at that time, but certainly over the years there's been quite a lot. That wasn't really addressed until the BLAST project, the one which we're just coming onto, which was a major project, again coordinated largely by Jack Dongarra, on a new BLAS, an updated BLAS, which would involve the legacy BLAS, which was level one, two, three that we had before, but would also include a mix-precision BLAS. It would include new BLAS which didn't exist for dense matrices, and also include the sparse BLAS. Indeed, I was working with sparse BLAS quite well before that.

Now, in direct methods for sparse matrices, sparse BLAS are not actually to be encouraged. Some people have used them. But it's not to be encouraged, because it keeps the sparsity. As I said earlier on, from 1981 on I knew that the best way of doing it was to actually concentrate on dense kernels for sparse matrices. So the sparse BLAS was better for iterative methods, really, which was not my main area of research at that time, although I have done a little bit since. The sparse BLAS I worked on quite a lot with some people from IBM in Italy: Michele Marrone, Giuseppe Radicati, Carlo Vittoli, that sort of people that we had, with a proposal in the TOMS paper on sparse BLAS. [Duff, I. S., Marrone, M., and Radicati, G. (1992). A proposal for user level sparse BLAS. Report RAL 92-087, Rutherford Appleton Laboratory. Report TR/PA/92/85, CERFACS, Toulouse].

HAIGH: I saw an early proposal for a sparse BLAS in *SIGNUM News* in 1987. So that was actually the same time as the general level three BLAS.

DUFF: That's right. The *SIGNUM News*, was that not just for level one sparse BLAS? I think it was.

HAIGH: I think it probably was. I have the paper here, so I can check.

DUFF: It was John Lewis, I think, that was involved with that. I'm almost 100% sure that it was a level one BLAS. So certainly I knew about that work.

HAIGH: Was a level one sparse BLAS actually produced from that 1987 proposal? Would it be useful?

DUFF: It didn't really catch on too much, no, neither the naming conventions nor the interface so much. Partly because there's a limit. The trouble with sparse BLAS originally was all the indirect addressing that you have—if you're doing a genuine sparse BLAS at level one, it's a problem. With that problem, there wasn't a great efficiency, so the vendors didn't really think it was a good thing to do and it never really caught on. So one thing we were interested in was to have a higher-level sparse BLAS, which is essentially level two or level three, and there was more scope for doing efficient things. Again, that proposal was never fully adopted. Some vendors, IBM because there were IBM people involved in the paper, Cray because they were very keen to implement everything to try and see how the machine worked, and a few others did some implementations of our original sparse BLAS (that's the one with the Italians), but it wasn't really universally available. So when we joined in on this new BLAST project, this got developed quite a lot further. Mike Heroux, Roldan Pozo and myself got very heavily involved in developing the sparse BLAS, partly based on this earlier work with myself and the Italians and partly based on some stuff that they were doing or planning to do in C and C++.

HAIGH: I saw a reference in the 1997 proposal that the main audience of the proposed level three sparse BLAS would be producers of library software. Presumably they would be the ones writing the routines to do things to sparse matrices, so if you could get the library routines shifted over then the application programmers would follow.

DUFF: Certainly it's just that the idea was that you would have these kernels which were very efficient, very easy to use; they encourage modularity of the code, which is a good thing, and of course are completely portable with vendor-supplied (hopefully) super-efficient versions. And then somebody writing library software would then use that as a kernel operation and get the efficiency. Now, I'm thinking again mainly that people doing iterative methods for sparse matrices would use these kernels, because there you're doing things like sparse matrix vector multiplication.

HAIGH: Those people didn't take to the system as much as you would hope that they would.

DUFF: I don't know. I think it's still an ongoing thing. I'm not sure. The jury is still out on it. In fact, the jury is a little bit out on the whole BLAST project for the new BLAS. I know that people including myself have used the extended BLAS of Jim Demmel and Sherry Li and others. But the package as a whole? I'm sure it's downloaded a lot, people refer to it somewhat, so I guess it is being used, but I'm not sure if I've got the full feedback on that. I guess we won't until Jack has written for another NSF grant on it and has to show how good the last lot was!

HAIGH: Would the sparse BLAS offer the same kinds of function calls that would be available to people using level three BLAS, but implement the internals a way that's optimized for sparse matrices? Or would the functions available be substantially different by those offered by the regular level three BLAS?

DUFF: The functions available are very similar. It's sparse matrix by vector multiplication, sparse matrix by full vector multiplication, although some people would like to do sparse matrix by sparse matrix multiplication. We also have sparse matrix by dense matrix multiplication, so that's the level three one. There's slightly different functionality. They're somewhat different, and somewhat more limited, I would say, in the sparse BLAS than in the dense BLAS. But really, we're trying to target operations that are used routinely in the iterative solvers.

HAIGH: Now, it seems that the motivations behind the original BLAS were to be able to produce portable high performance code with a good degree of modularity, so that the code would also be maintainable and readable. Were those basically the same motivations that have continued to drive it, in a sense?

DUFF: Yes. Both the BLAST project and the sparse BLAS.

HAIGH: I haven't heard as much in my other interviews about the BLAST project. Is that essentially a "level four" BLAS – a higher level of abstraction?

DUFF: No. BLAST is…well, it's all levels. It stands for BLAS Technical forum, but just the T I think has remained. There's a website hosted by the University of Tennessee on this project. It has the software available and some model implementations of the software. As I said, it's not quite the same status as the older BLAS—BLAS-1, BLAS-2, and BLAS-3. I'm not entirely clear as to the take-off with vendors and users in general. People do use it, and I've seen reference to people using it. As I said, I myself have just recently, in some work with a student I mentioned that's in California just now looked at using both the mixed and the higher-precision BLAS, which were included in the BLAST project. So the BLAST project was not level four or anything; it was just revamping to some extent what was there before, rebadging some of it, and adding a few extra functionalities which wasn't in it.

HAIGH: So it's not building up on any bigger level of abstraction; it's just tidying up and improving what's already being done by the time you get to level three?

DUFF: All levels, yes, because it's got the whole legacy BLAS there with one and two, so that's right. I think most of the additional BLAS that I did were the level three ones, with maybe the odd level two. Tidying up things like complex arithmetic, trying to get things uniform throughout.

HAIGH: I also saw references to a RISC BLAS.

DUFF: That's right. That was more implementation-wise. That was with Michel Daydé and others in France. The various architectures of machines and the various names and things that happened over the years, and there was a while when the so-called RISC, or Reduced Instruction Set Computer, processors were in vogue and everyone wanted to be efficient on them. So we were writing BLAS that was efficient on RISC processors. But the functionality and the interface was just the same as the BLAS.

HAIGH: And you did some work also on implementing it for various kinds of supercomputers?

DUFF: Yes. That's right.

HAIGH: So people have talked about the difficulties of getting things to work on vector machines and of course on parallel machines. What was it about RISC machines that made it hard to get high performance without rewriting the BLAS?

DUFF: I think it was just that the compilers weren't able to take the straight FORTRAN code, for example, from the BLAS and implement them efficiently, so we had to do things like loop unrolling and this kind of stuff explicitly in the code to do that. I think the compilers caught up a little bit later. I don't know if we had waited a bit, they might have been able to do a reasonably good implementation.

HAIGH: That was the theory behind RISC in general, right, that you can put more intelligence into the compiler and have the chip itself be simpler?

DUFF: That was the idea. The rest of the chip was very simple; that's right; the hardware was very simple. But because of that, you had to do things like loop unolling to get the instruction set feeding in at the right pace.

HAIGH: So it seems like the only remaining thing we have to talk about is your general involvement with the international supercomputing community. You've talked about aspects of that already; you've talked about writing your report for the European atomic energy community and about work for the United States and Hong Kong governments. But I know that you've been involved in other countries as well. So does this all go back to your experiences visiting Argonne?

DUFF: Yes, I guess that was very much the seed corn for it, because when I was at Argonne it was just the time when a lot of different machines were coming out. I had this report that I did with Jack Dongarra looking at all the machines that were around and writing out a coherent description of each in the same sort of format. As I said, that formed the background to a congressional report, so I was very pleased to be up on The Hill, as it were, with this data. So I guess that was my interest, and presumably I got known as well from that, and from the one I was writing for the European Union. I was very much involved with the Cray 2 at Harwell, writing papers associated with algorithms there and I gave quite a few talks in the Cray User Group. So one thing or another, one's name gets known, and people say, "Right. Who should we have to review supercomputing in this country?" And it's usually a small committee, sometimes a committee of one, and you're invited to do things. So I was invited to do things.

One of the first places was Malaysia. I had not been to Malaysia before, and I like traveling. So I went along and had a fairly intensive time. I was doing a review for supercomputing. It was very easy, because there wasn't much supercomputing there at that time, and their requirements were not quite as sophisticated at that time as perhaps they were here. I came back later on to investigate the mathematics department and make suggestions for mathematics courses. This is something that any professor gets involved in to some level or other. They want to know whether they're working at the level of Western Europe or the USA or whatever, so they invite you to look over all the courses, meet some students and talk to the faculty, even to evaluate the course structure.

The Swedish one was a bit different. I guess I knew people like Bo Kågström. He now still runs the big supercomputer center in northern Sweden. So we knew him quite well, and I knew several other people in the Swedish supercomputers scene. So I was invited in a subcommittee which involved Bill Buzbee, Ralph Roskies, and Jack Dongarra, to go around to all the supercomputer centers in Sweden—there were about four or five of them at the time—and to write up comments. In fact, we had been asked which of these should be supported by the Swedish government. So you can imagine that we were getting lobbied quite a lot both by our friends that we knew and by other people. We actually did have quite an influence, I think, because at that time far and away the main government-subsidized supercomputer was Linköping in mid to southern Sweden, not far from Stockholm, with the Saab company. The thing that we thought really could be improved was that Saab had control of this machine. SAAB did a lot of defense-related things—really more SAAB aircraft than the cars that you know about—and because of that the academic community were not actually prime users of the machine. So government money was going there and was being used by Saab for their own computing. I'm sure it was all properly accounted for. But the supercomputing community didn't seem to do as well, and there was sometimes down time because of the security, et cetera, et cetera. So what happened was there were some changes made in Linköping after that. It continued to be a supercomputer center, but the relationship with Saab was changed I think for the benefit of the academics. And some of the other centers were boosted by our visit. In particular, PDC in Stockholm, did a terrific presentation, one must say, and they did well out of that, and got some more money to develop theirs. The people in the North, Bo Kågström that I mentioned, we gave them such a good write-up and said that they should get money, so they actually got money from other foundations in Sweden, perhaps on the back of our report. So I think the Swedish supercomputing was helped by our visit and definitely got developed.

So these are kind of exciting things to be involved in. You're not exactly doing research, and they don't get you Nobel prizes (not that we don't get that much in maths anyway) but, really, they do have a high effect on peoples' lives and on research, so that's good.

So after Sweden, I was asked to review the Netherlands. I guess this was my friend Henk van der Vort that put them in touch with me. That was a smaller committee. I can't even remember who else was on it, but it wasn't the same people as before. Again, it was quite arduous, and we had to interview many people. It was more for the users of supercomputers that we were reviewing and what was the requirements rather than for the sites for the supercomputing.

I should say, in between times I was heavily involved (I forgot to mention it and it doesn't really appear on my CV) as chairman of a steering committee for a benchmark working group in the UK for the purchase of supercomputer to replace the Cray YMP, which was at that time here at Rutherford. But that was split into two parts. One was to choose the site, and the other was to choose the vendor. It was a peculiar system that they had there. We came down in favor of a Cray T3E against several other companies. I probably shouldn't be saying who they all are because that's semi-secret. But I should mention that KSR nearly threatened to sue me because they claimed that I had discriminated against them, until I pointed out that it was their own machine that discriminated against them. So it was interesting. You have to base your judgment on of benchmarking, which we had. We had all that data. So that was fine. But that was a really very interesting and very harrowing experience, really. It was really like being cross-examined. They were trying to make millions of pounds selling machines, and if they think that you're not giving them a fair crack of the whip, they really go for you. So I had to defend myself and my committee, which was quite harrowing. But we did it successfully because we were scientists and doing a good scientific job.

[Tape 5, Side A].

Session Four concludes.

DUFF: So sometimes, other countries were easier. Of course, I was very much heavily involved in various other technical option groups, like I mentioned in Britain for computing, mainly just before I left Harwell. The most recent one was being involved with the Helmholtz Institute. Helmholtz is a centrally funded number of labs in the Federal Republic of Germany, which is a fairly new amalgamation. We had to review that process, and Tom Manteuffel, who was the President of SIAM, was chairing that committee. I was involved, and it was a very big committee with lots of quite distinguished people on it. I don't claim to be such, but there were a lot of other people who were very distinguished. That was very interesting, and ran along quite efficient lines, as you can imagine in Germany. I guess our recommendations were taken on board there again.

Just about the week after, I was in Portugal, which was looking at courses and CMAP and the structure of the mathematics department in Porto. Along the lines of what I'd done in Malaysia. Both pure and applied math, I had to sit through a couple of pure lectures which I couldn't understand much of, but none of the people doing pure maths could understand them either, so I felt happy about that.

HAIGH: It looks like you had done something in the Netherlands as well.

DUFF: Yes. I sort of half-mentioned that. It was Henk Van der Vorst who put me in touch with that. That was more the users we were talking to about what the requirements would be and what type of machine they should get based on scientific requirements of the users. They got quite involved in that. That was pretty exhausting, and there was quite a lot of science in that actually.

HAIGH: So I'll just hand you back this portion of your resume and you can tell me if you think there's anything important that we have missed. That's professional activities and memberships. Then there's the various editorial boards, journals you've reviewed for, and theses examined. You mentioned a number of the students. Are there any students you've worked with there who you think you should say more about?

DUFF: This is mostly people I've examined, so in many cases I have not been the supervisor. But in France it's not uncommon to have the supervisor in the jury, so there are a few there. I think I've done more than my fair share of having to examine students, and I've quite enjoyed it in many cases. It does require a few hours or a few days, really, to read manuscripts and that type of thing. Obviously professors give you theses to examine that they know that you're interested in and are in your area, so it's almost a bit better than reviewing papers that you might get from editors. I quite like examining theses, but I like to do no more than a couple a year, if I can get away with it.

Now, looking at this list is triggering some other comments about current research, but I'll come onto that. One of the things, in fact, that we got involved in through one of the students I noticed here, Bruno Carpentieri and subsequent students at CERFACS is in electromagnetic radiation, which is something that I had not done anything with at all before. That's interesting in the sense that the formulation that we were using, boundary-element methods, doesn't give rise to a sparse matrix. It gives rise to a dense matrix, in fact. But you can use sparse techniques in the pre-conditioning phase to solve this by an iterative method. So that brought me into the idea of pre-conditioning, which is quite a big area of Luc Giraud, who is my second-in-command at CERFACS and has just left to become a professor in the university. I got collaborating with him over this, developing sparse techniques for the dense problems from boundary element methods, and working with fast multipole methods and other such things. That was quite a big area of research that has only been started, really. Looking at the dates of these people, it was only started about 1999, 2000. But it's still continuing somewhat. We had a student who just finished this year who had done a thesis on that.

HAIGH: Anything else about the students supervised?

DUFF: Most of the ones in the past have gone on to be professors or academics.

HAIGH: So are there any other aspects of your current research and interests or future plans that you would like to talk about?

DUFF: Well, one of the duties I have coming up in 2008 is this very famous research assessment exercise, RAE. In the US, this probably doesn't mean much, but in the UK, part of the government money for universities comes from teaching and part comes from research. They're both reviewed, in different ways. The research assessment exercise that I'm on is reviewing the research capacity/potential of universities and assigns a numerical rating to the department. This numerical rating will party define their funding. So it's really very important for universities, and they get quite excited about this. This exercise takes place not in quite fixed intervals, around about every six years or so. It will be 2008 before we actually conduct the review of this, although there will be some preparation discussions in 2007. But already we're preparing the groundwork and the rules by which the universities will submit their entries. We're revising and collaborating with the universities to get rules that seem reasonable and they're happy with. So that's a major undertaking, and there are only 15 people to cover the whole of applied mathematics in the United Kingdom, including Scotland as well. So in that way can I expect quite a lot of work. Somebody said it will be equivalent to reviewing 200 papers each.

HAIGH: Is that the kind of thing that has the potential to make you unpopular?

DUFF: Well, I guess. Whenever you're talking about money and people getting money or not getting money, I suppose that's it. I would hope that my personality and my previous interaction with people, the fact that I hope people regard me as being a fair and reasonable person, nothing personally should come back on it, I would rather hope. But yes, it's got the potential to make you unpopular, sure, I guess, and the potential to make you extremely popular as well with people who do well.

HAIGH: Any other aspects of your current work you think you haven't dealt with already?

DUFF: There is really the development of how the group at Rutherford has developed and how the group at CERFACS will develop, and how I'll personally develop. I'm not so far off 60, and when I get to that age my contract should expire here, but of course both the United Kingdom and the European laws are such that that wouldn't really be enforceable if I wanted not to go. In fact, it would be unlikely that they could force anything before 65. But having said that, it's attractive to me perhaps to not have so many administrative duties, so I'm quite looking forward to maybe ceasing to become leader of numerical analysis here, and (perhaps not at the same time but shortly afterwards) trying to hand the reins over in CERFACS to somebody else. I'd still maintain some kind of links, spend more time doing research, in fact, after that. So that's the sort of plans I've got personally, which could have an implication for the two groups and the structure of what happens there as well.

It should be said that over the years, the group at Rutherford has established itself as being the prime place for large-scale numerical analysis research into large-scale computing in the UK without any shadow of a doubt, and that's why we get our grants and why we get our money. Of course, the intention would be that we continue even if my role might change a little. The other people who are here are more than adequate for carrying on that role, and hopefully we should be able to maintain that. But it's always a dodgy business. It's fairly soft money. We have to apply every four years to get our grant renewed. The next renewal is in 2007, so that should be interesting.

So there's a lot of slightly imponderable issues ahead. Research-wise, there's still mileage in multifrontal stuff. The MUMPS package is still being developed, although I'm not doing any of the direct coding of the MUMPS package anymore. And quite a lot of research has involved students of mine and students of other people. So it's really quite a major project. In a sense, Patrick Amestoy and Jean-Yves L'Excellent are coordinating it almost more than me. So the good thing there is that if I do take a lower profile, it's not going to destroy that project; it will continue.

I seem to spend a lot of my time recruiting people as well. We are making recruitment right now at Rutherford. Fortunately and happily, because I am trying to change my role here in the future, other people in the group are taking the lead role in this recruitment. So that's excellent. At CERFACS it's just ongoing because we have a lot of people only on two-year contracts, so we're always employing new people. So it's very exciting in some ways, but quite exhausting in others.

The other thing I would like to do—it's quite good to record it now, because there might be a resolution by the time anybody listens to this—is to try and find somebody to take over the *IMA Journal of Numerical Analysis*. We didn't really talk about that too much. I've been Chief Editor originally with Alistair Watson since I think it was 1991 or something. That's been a very major thing. It's one of the top general numerical analysis journals, in the world, and it's run under the IMA auspices. Alistair and I were Co-Chief Editors, then he, who is slightly older than me, retired, and Endre Süli from Oxford took over as the other editor. I was rather hoping to persuade some colleagues to take over from me towards the end of this year or next year, but at the moment I'm not succeeding in getting somebody. But hopefully we'll find somebody within the next year and I can pass that on. It certainly would be my target not to be the editor of the *IMA Journal of Numerical Analysis* by the time the RAE clicks in fully, because I do spend at least half a day a week on IMA Journal business.

Nothing else for the future. I don't know really. I'll be dead eventually, but a lot of things to do before then.

HAIGH: I'll ask you the standard two questions that I have asked everybody at the end then. The first one of those would be, looking back over your career as a whole, what would you say your biggest regret is, either in terms of something that you did or just in terms of something that you think was a good idea but the world didn't respond to in the way that you might have hoped?

DUFF: It would be terrible to say I don't have regrets. There are little points at life in which I had decisions to make, like, should I have left academia and Newcastle and come to Harwell? Should I have taken the post at Strathclyde when it was offered? Should I have migrated to the USA? These are all kind of major decisions, sometimes were made in a short period of time, sometimes over many months or years. I can't say that I regret now, looking back particularly at any of then. It would have completely changed my life and my family's life, and everything else totally. But the main thing is that I'm happy in the work and I'm very happy with the people I've met at work, because I think that our community is one of the best communities for research activities. I look at physicists, chemists, and see them cutting each other's throats and stabbing each other in the back. Not all of them, I should say. That's a bit rude. But a little bit more of that goes on than in our community, which is pretty good. Nothing perfect, but pretty good. Things like the Householder Meetings are really very nice. It's a pity that they're selective in a way, because some people that are not invited do feel excluded. But they keep the numbers down, so that you have to select. But having said that, the atmosphere is nice and the community is very good. That's the happy side of things, if you like.

Regrets? Gosh, it sounds very smug or that. It's not that I've never done any mistakes, really. It's difficult to say. With my organizational interest, maybe I should have gone back into universities. One of the major universities in Britain offered me a professorship (I wouldn't say it was in concrete in writing, and I shouldn't mentioned which university) just before the most recent research assessment round, when of course my research publication record would have done a great deal of good for the department. That's always a thought. If you go into this, there's a career or structure above being a professor in a university and being Vice Chancellor, or being involved in the political side of higher education. The universities definitely have an easier route that way. In the quasi-civil service we're in, I don't have that possibility. So when my career comes to an end, it comes to an end. I'm hoping that I'll do more research and less administration, but I wouldn't get knighthoods. If you're in universities you can go on to become far more famous, as it were. I have chosen not to do that. I don't think it's a regret, but it's always something that you wonder about.

HAIGH: Then the more positive reverse of that, what do you think is the single achievement of your career that you're most proud of?

DUFF: One thing I'm very proud of, which is not just a research item, is really getting CERFACS up and going—my group at CERFACS, really I should say. I don't want to take credit for the whole place. Pierre-Henri Cros is much more important in that respect. But I think it's a fantastic group. I think we have actually changed my part of French science. When we started CERFACS, numerical linear algebra was almost non-existent in France (not quite, it would be rude to some people). And now it's really becoming quite a big topic, and a good lot of the people are people who went through CERFACS, had been influenced by people at

CERFACS, or whom because of my French connections got involved at various levels when I have been there. So that whole development, that whole area of that kind of the slightly more experimental side of science, which was never there in France in the early days, is now there. Some people do recognize this. I'm not sure that you get full recognition, because people don't like to recognize the contribution of foreigners, particularly when they come from Great Britain. But that, to my mind, is one thing I feel really happy and proud of. If I was lying on my death bed thinking what use I had been to the world, I would think of that.

I am also extremely happy that I managed to keep this group together that's now at Rutherford when we had absolutely been faced with complete destruction at Harwell when they were changing their whole ethos on research. So I'm very proud and happy that we survived and still flourish. So these are not really research things.

The research things, I guess, are things like this whole multifrontal bandwagon, if you want to call it that. So it was myself and John Reid that had the initial ideas after this Speelpenning work. Really, I processed that not on my own, but because of people like Patrick Amestoy and others who were helping me, but we've really developed that enormously. It's one of the main ways of solving sparse matrices. There's no question.

I think there are a lot of other very solid contributions. Some of the work on the combinatorial side has been important, and it's coming back again, which is nice. So this work I did say in 1979, '80 is coming back now and being used by people nowadays. Not that it was entirely unused before, but it's coming back a bit stronger. And some of the recent work, a very simple idea that I did with a student of mine, Jacko Koster on this permuting large entries to the diagonal has had a very strong influence way beyond the amount of effort that we put into it. It had a strong influence on preprocessing systems for both direct and iterative solvers. So I'm very pleased with that work as well.

I'm just scanning over my ex-students to see if there's something special. Sparse matrix work is fun. I wouldn't say that we have completely pushed the envelope of the field anyway, but we have done good work there and it's recognized. Again, one of the positive things is working with the younger people and seeing the subject develop. I wouldn't have had the chance to do that too much at Rutherford. We do have CASE students who are sponsored by places like Rutherford and. I have supervised one of these, and other people in the group have done that as well. But at CERFACS we've got four positions for PhDs at any one time, and usually, I'm supervising at least one and possibly two at any one time. So that's just great, this interaction we have. The international flavor, of course, from CERFACS is very good.

The things like the BLAS are very important, but these are big efforts. My part is certainly no more than the quarter or less of the total. I don't feel as if I was big driving force there, although I was probably using level three BLAS before almost anybody else. So that paper in 1981, at Dundee that I mentioned to you earlier, "Full matrix techniques in sparse Gaussian elimination," I think that was one paper I still feel very proud of, as I do for the old IEEE Proceedings paper, but that was a survey. Research-wise, I think the Dundee paper was good. But it got a little buried in the conference proceedings.

HAIGH: That concludes the questions that I have prepared. If you have anything else to say during the interview, then now would be the time.

DUFF: I think it's been an interesting exercise having the interview, and I'm sure that when I have seen the transcript it will trigger other things. Obviously, if they're not too major, I may think of adding them or something. But I actually have got a project to organize my room a bit more. You see there are quite a lot of papers around. No doubt, I'll unearth other things there. So I may send in a supplementary tape or comment to you at a later date. It's been quite fascinating looking back over the years to see what's happened, and I've been made aware by your questions and preparation that there's more to my life than I thought, I guess. So it's all been very interesting, and it makes me feel extremely old.

HAIGH: Thank you very much, then, for agreeing to take part. We will conclude things there.