An Interview with

JOHN PINKERTON

OH 149

Conducted by John Pinkerton

on

23 August 1988

Esher, England

Charles Babbage Institute
The Center for the History of Information Processing
University of Minnesota, Minneapolis

John Pinkerton Interview
23 August 1988

Abstract

Pinkerton begins by discussing his education and wartime work in radar technology in England. He then describes his movement into the computer industry after World War II and his work on the LEO I and LEO II computers. In this context he discusses the British computer firms J. Lyons and Company, Leo Computers, English Electric Co., and International Computers Ltd.

JOHN PINKERTON INTERVIEW

DATE: 23 August 1988 INTERVIEWER: John Pinkerton

LOCATION: Esher, England

This is John Pinkerton. I am recording this tape in my house at 2 Sandowne Road, Esher on the 23rd of August, 1988. I am going to start this first tape with a certain amount of chronology, which may be useful in editing the rest of the recording on this and other tapes.

I was at King Edward's School in Bath, which is a grammar school, from September, 1927, to July, 1932. I was at Clifton College, Bristol, as a boarder, from September, 1932, to July, 1937. And then I was an undergraduate at Trinity College, Cambridge, reading for the Natural Sciences Tripos from October, 1937, to June, 1940, by which time the war had been going on for about ten months. I then went to work for the government at the radar establishment in Swanage and later in Malvern from June, 1940, to August, 1945. I returned to Cambridge in October, 1945, as a Ph.D. student in the Cavendish Laboratory and I remained there until the end of December, 1948. In the end of December, 1948, I got married, and I started work with J. Lyons and Company, Ltd. in the middle of January, 1949, with the intention of building a computer for them. I remained working for Lyons directly or indirectly until the merger of Leo Computers, Ltd. with the computer department of the English Electric Company in 1963. I then was transferred with the English Electric Computer activity to form ICL by a merger with ICT in 1968. I finally retired from ICL in August, 1984.

To get further details of the chronology during the Leo period, you can consult Mr. Thompson's diary, which goes from the date they first thought of having a computer at Lyons to about 1963 or 1964.

I'm now going to try and answer the questions in Dr. Aspray's letter to me, which I have numbered. The first question was: Science quarters in the laboratories at Clifton, what I learned there?. The first thing to say was that Clifton had built for itself, as a result of a very generous gifts from rich, old Cliftonians, an extremely magnificent set of science laboratories which were finished in the 1920s. There were, if I remember, two very large physics labs and one smaller physics lab; two very large chemistry labs, and one smaller chemistry lab; and two physics lecture rooms and two chemistry lecture rooms. In addition, there were the usual offices for the staff and there were preparation rooms for the science assistants to get experiments ready, and so on. It was all extremely well organized. The head science master, when I was there, was a Mr. E. J. Holmyard, who wrote a great many chemistry textbooks and was extremely well known for these books at the time. The chief physics master was a Mr. W. C. Babcock. There was a total staff of about 8 science masters altogether.

The courses were geared to the examinations which boys at the school regularly took. These were known as the School Certificate and the Higher School Certificate. Today, these have been replaced by what are known, or have been until recently, known as the Ordinary Level and the Advanced Level. But I don't know what the corresponding exams would be called in the United States. The advanced level is the sort of level of attainment that a student would have when he went up to university and, I think, is rather more advanced than is customary for the first year students in American universities. But I'm not sure about this.

The subjects were chemistry, as a separate subject; physics, as a separate subject; and biology. Most boys did not take biology, but the biology course was designed for people who were going to go into medicine and wanted to take their first medical examination from school, which you could do. Which, in fact, I did because my parents thought it would be a good idea if I could become qualified as a doctor. This was something I did not wish to do and, in fact, I eventually failed to get the first medical examination at London University as an external student. However, by that time I had got an open scholarship to Trinity College in Cambridge and I didn't feel any great sense of failure at not being able to get my London first M.B exam.

By the time I left Clifton, I had a very good grounding in chemistry, both inorganic and organic chemistry, and in physics as well as in biology. I also had a reasonable level of advanced mathematics. This was, in fact, something of a pity because it meant that I was a little bit idle in my first year at Cambridge because I found that much of the work that I was expected to do then I'd already done at school.

The second question Dr. Aspray asks is: How I came to choose Cambridge.

The answer to that is quite simple. The school recommended that I should try for a scholarship examination at Cambridge. At that time, you had to pick one of a group of colleges as you could not apply for scholarship exams to all of the colleges, simultaneously. And so it was a bit of a lottery. I was advised to try the group with Trinity College in it, because Trinity College was regarded as the best college for natural science students at the time. And in fact I was awarded an open scholarship by Trinity as a result of an exam, which I took in the latter part of 1936. You then waited at school until you were old enough to go up to college in the following October. In addition to getting an open scholarship to Trinity College, I also got a so-called state scholarship, which was awarded at the time on the results of the higher certificate examination. In fact, I never got any money from this because my parents' income was too large for me to qualify; but had either of them become ill or died, then I could have reapplied. So it was regarded as a sort of insurance policy. I was also assured that if I didn't claim the money, it would not be totally wasted but would go to other candidates who were lower down on the scale.

The next question that Dr. Aspray asks is: What I learned from my Cambridge tutors.

I think the terminology at Cambridge is a bit different from that at the American universities and I may have to explain how the college teaching and the university teaching fit together at Cambridge because the system there, and at Oxford, is rather different from that at almost every other university in the United Kingdom and probably different from that at any other university in the world. The point about it is that the college makes itself responsible for seeing that you entered the correct courses, which were suitable for your attainments and for the profession you wished to pursue when you'd finally taken your degree. So the college had what they called directors of studies, who merely advised you which course to take. There was also college tuition, which was intended to supplement the university tuition and helped you to sort out any problems which your lecturers didn't seem to be explaining properly. The people who did this were known as supervisors who were college dons who were knowledgeable about particular subjects. I can remember one or two of the people from whom I had supervision, who were in fact quite famous people. There was Charles Coulson, who was a mathematician. There was Maurice Price, who was a physicist, and, I think, held university appointments in California as well as at Bristol and Oxford in the United Kingdom. There was a physiologist called Routin (?). I think I also had supervision from Norman Feather, the atomic physicist. I can't now remember the names of all the other supervisors, because you could have had a new supervisor each term in the different subject. So quite a number of people were doing this. I've tried to remember, also, the names of some of the lecturers whose lectures I went to. One of the most successful lecturers in Physics was Mr. J. A. Radcliffe, the ionospheric physicist who was extremely good at electromagnetism. I also went to lectures by the venerable Alex Wood on acoustics. I seem to remember some lectures, or at any rate a practical class being organized by the very venerable physicist called G.F.C. Searle, about whom there were many stories. And I certainly remember lectures by Sir Lawrence Bragg, who wasn't, incidently, a particularly good lecturer. At least I didn't think he was. In the first year of the war, the physics part two course, which I attended, was a little adapted to include a certain amount of electronics. One of the speakers on this course, who spoke about electronic circuit design, was Sturley. In the part one of the Natural Sciences Tripos you were required to read a total of three subjects. As well as being whole subjects there were half subjects. The subjects I actually took were chemistry, which was a whole subject including both organic and inorganic chemistry; physics; and half-subject mathematics with half-subject biochemistry. Having taken quite an interest in biology at school, I found biochemistry a stimulating and interesting subject. When it came to the so-called part two of the Tripos, you were expected to specialize in a single subject. I had some difficulty, initially, in deciding whether to read chemistry or physics. But eventually, since I found I was better at practical physics than practical chemistry and since I came to the conclusion that the practical side would be what mattered when you eventually got a job, I opted for practical physics.

The story of how I got into radar, I think, is worth recording. It was a very strange story. I had been interested in radio, in one way or another, since a very early stage. In fact, since my father built a crystal set in about 1923 and I started listening to the children's hour. I can just about remember this crystal set and listening to it on headphones. Then, when I was about seven or eight, somebody gave me a book of circuits which I didn't understand; but I spent hours poring over these circuits which included the old fashioned bright-emitter valves, each one of which had its own heater rheostat which you adjusted so the brightness of the filament looked about right. That was the only way you could do it. By about the age of somewhere around eleven I started to build radio sets, which actually worked --

starting with a one-valve set and gradually adding numbers of valves until eventually I was building six and eight valve superhets with some sort of pretensions to high fidelity as understood at the time.

When I went up to Cambridge, I had my own superhet receiver with a push/pull output stage and I used to invite friends around to listen to gramaphone records on this machine, which gave a reasonably satisfactory sort of reproduction. I also joined the Cambridge University Wireless Society and had some hope of becoming a radio transmitter. To do this, you had to pass an examination in the Morse Code. You had to be able to send and receive Morse Code at twelve words per minute. I never succeeded in doing this so I did not, in fact, obtain an amateur license. But it was quite important eventually, as it turned out, for my career, because through the Cambridge University Wireless Society I had met Maurice Wilkes before the war when he was a research student; that was one reason why I got into computers much later.

However, at the start of the war, or before the war started rather, it was customary for people who were going to read physics to go to Cambridge in the long vacation, that was during July and August, to do practical work in physics because the number of experiments you were expected to do in the course of your last year was too many to get fitted in in the ordinary terms, which only lasted eight weeks anyway. During this period, a mysterious notice went up on the board in the Cavendish Laboratory, signed by Mr. Radcliffe, to the effect that there was going to be a special government course for people who had a particular interest in radio and anyone who wanted to go on this specially organized course should be in touch with him. This I did and he was very mysterious, indeed, about it, saying that I would have to be cleared by security, whatever that meant. Well I didn't know what it meant, but two or three weeks later he said this had been done and that I would receive instructions from him as to where I was to report myself. Well, in due course, and as it turned out exactly a week before the war eventually started, I was summoned to meet him and several other people from Cambridge at a hotel in Folkestone. There we were initiated into the mysteries of radar, and we were told that the course was a crash course to learn all about radar. In fact, it started almost the day the war itself started or a few days beforehand at the radar station on the top of the cliffs at Dover. I remember hearing Mr. Chamberlain's broadcast to the effect that at 11:00 on Sunday, September the 3rd, we were at war with Germany. We all sat there and calculated how long it would be before the German bombers came over and destroyed the radar station, which we assumed they would do in order to be able to initiate bombing raids on Britain. Of course, as everybody now knows, this did not happen; but we thought at the time it would.

When the radar course at Dover was finished -- it took about five weeks and we felt we had got a fairly good appreciation of the way the radar systems worked and were designed, I was offered the opportunity of joining the RAF with an immediate commission. However, having heard that there was a research establishment at which radar systems were being investigated and designed, I said I would much prefer to join a research team; and after some interviews with the University recruiting board at Cambridge had been set up to direct graduates or prospective graduates into suitable war jobs, I was told I could finish my degree. So I went back to Cambridge in October, 1939, and finished the degree in the summer of 1940. I was posted to the government radar establishment in Swanage, where I did a number of very interesting things. Almost the first job I was involved with was setting up one of a chain of listening stations to listen to the long-range Lorenz navigation beams, which the Germans were then using for directing their bombers over this country to drop bombs on places like Coventry. I was involved with setting up listening posts to listen to these signals, first at Scarborough and then near Dundee in Scotland. For the first year or two, I was concerned with aerial impedance and polar diagram measurements. Then, later on, I got involved with the design of an airborne interrogator for coupling with various airborne radars to receive IFF mark III signals, which were used to distinguish friendly aircraft from hostile aircraft, particularly in other airplanes fitted with airborne radar systems. (IFF = Identification Friend or Foe.)

Anyway, by the end of the war I felt I knew quite a lot about radar and quite a lot about pulse circuitry and I was very fortunate in being offered the opportunity of going back to Cambridge as a research student. I had been awarded what was called a Senior Scholarship by my college, and, also, a grant from what was then known as the Department of Science and Industrial Research, DSIR, which Mr. Radcliffe arranged for me. So I had both a DSIR grant and a college senior scholarship, which were not particularly riches but it was enough to keep a research student going. Initially my idea was to use ultrasonics to build equipment for exploring the insides of living creatures, because I had seen a piece of ultrasonic equipment at the radar establishment in Malvern. This was in fact a trainer for the airborne radar H2S designed by Dr. Uttley. It simulated the effect of flying over the ground with an H2S radar by using a tank of water and an ultrasonic beam with a model of the terrain on the bottom of the tank. Because the velocity of sound and the velocity of electromagnetic waves are very different, it is possible to make a scale model in a tank of reasonable size. I think this trainer had a tank of something like eight or ten feet square. This gave me the idea that perhaps the ultrasonic signals could be made to pass inside living creatures, say man or any other animal, and show up what was inside. However, Mr. Radcliffe's advice was that this was more in the nature of an invention than a research project and there was no real guarantee that anything much would come of it. In fact, I think this was very good advice because at the time electronic techniques, particularly techniques of presenting information on cathode ray tubes were fairly primitive; and had I pursued it I probably wouldn't have made it work and I probably would not have got a Ph.D. out of it. Mr. Radcliffe, being a director of research, was more concerned that I should pursue a line of inquiry that would lead to a Ph.D.

I was a bit of an oddball in the radio section of the Cavendish because I wasn't strictly doing radio, I was doing ultrasonics. What I eventually decided to do was to investigate the anomalous or apparently anomalous absorption of ultrasonic waves in ordinary liquids, and try to build apparatus to measure absorption very accurately. Previous measurements had been rather inaccurate and they weren't very extensive. So there was some doubt even as to whether absorption was strictly anomalous or not. But this allowed me to exploit pulse techniques and I built some fairly elaborate apparatus which has all been described in literatures, so there's no point in my going into that. With hindsight, the benefit of it was that the techniques that I used were exactly those required to build a store for a high-speed computer. I didn't know that at the time, of course, it was purely coincidental. The other pure coincidence was, as I said earlier, that I had come to know Maurice Wilkes and I had seen him occasionally during the war. He wasn't a close friend. He was more, I would say, an acquaintance. However, he was building a computer and the computer was going to use ultrasonic storage. So I was naturally taking a mild interest in this and went to various lectures and colloquia that he gave about his plans. So I was generally aware of what he was doing. I also heard from him during the summer of 1948, that Lyons were interested in building a computer, which he thought was rather a curious idea altogether. Indeed it was a rather curious idea that a catering company with no background in ultrasonics or even in electronics should contemplate building anything as advanced as a computer. I didn't know, of course, when I first heard about this, that I was actually going to be involved in it.

During my last year at Cambridge, I was looking at a great number of jobs, and I think I must have been interviewed at fifteen or twenty different laboratories up and down the country to see whether they were doing the sort of work I would like to be engaged in. But I had not accepted any offers until towards the end of the first academic term in the academic year starting October, 1948. I had tried to get a college fellowship at Trinity College, but I was unsuccessful in the fellowship competition. So I was advised, again by Mr. Radcliffe, that I really should look for a career in some industrial laboratory. Towards the end of that term, I saw an advertisement which was placed by J. Lyons, in the magazine *Nature* I think. They didn't quote Lyons, but from what I'd heard from Maurice Wilkes I suspected it might be Lyons who were advertising and indeed it was so. The story of how I got this job has already been recorded; it's in the paper in the *Annals* so I don't think I need to repeat it again.

One of the questions you ask: what was the relevant technology with which I was familiar when I went to Lyons.

Well, I suppose I can summarize it something like this. As far as circuit technology went, I was familiar with the usual designs of time-base circuits, of audio frequency, video frequency, and RF frequency amplifiers including wide band amplifiers, design of simple oscillators, dividing circuits, and the design of CRT displays, because all of these things figured in my ultrasonic research apparatus. In addition, I was familiar with ultrasonic propagation and the whole physics of this business which entered into the design of my research apparatus. I was familiar with the design of driving circuits for connecting quartz crystals to oscillators. And I was familiar with the physics underlying the design of ultrasonic delay lines. So I really was in a pretty good position to undertake the detailed design of the circuits of a computer. I could also say that I was familiar with these things not only from the theoretical and design point of view, but from the point of view of actually building them and making them work. Also I had had some experience of supervising contracts with several contractors during the war who were making the interrogators for airborne radar in conjunction with IFF Mark III. There were at least four of these contractors. So I was used to writing the government specifications for the equipment and also seeing to it that the manufacturers actually produced equipment that conformed to the specifications. So I was used to the process of testing prototypes and getting the Aeronautical Inspection Department, represented by government inspectors in the factories, to see that all the production models came up to the standards they had done in the specification. All of this experience was, of course, very valuable when we came to put contracts out from Lyons to get the parts of Leo Mark 1 and Mark 2, manufactured by somebody else.

You have a question about interaction with the NPL.

I would say that, broadly speaking, once we knew what was happening at the NPL, what we did had practically no impact on what they did and what they did had practically no impact on what we did. This was because we were following the design of EDSAC in building Leo and only departing from it when we felt we had to do so. We established the philosophy at the beginning that we wouldn't change anything unless we knew exactly the reasons why it was designed in the way it was in the first place -- which, looking back on it, seems to have been a very sensible philosophy. We knew, for instance, that the NPL were adopting optimum programming devised by Alan Turing, where you arrange matters so that the next operand or the next instruction was just about to emerge from the delay line at the moment when you wanted it. This obviously made the whole system run several times faster than it would have done had you had to wait for it to emerge at a random time. However, this was not the way in which Maurice Wilkes had designed EDSAC, so it was not the way that we designed the LEO I.

Another question in Dr. Aspray's letter is: what knowledge did we have of the war-time codebreaking activities at RRE at Malvern.

Well, as far as I was concerned, no knowledge whatsoever. I don't think we knew that codebreaking was going on. We knew, at least several of us knew, that there was some kind of an establishment at Bletchley and found it very intriguing to try to find out what went on there. I didn't find out anything, whatsoever, until after the war was over. And even then I didn't find out very much. Despite the fact that my sister-in-law was actually employed there, I still got no information about it. We did, however, get quite a lot of the intelligence at RRE, which obviously had been derived in part from codebreaking activities but a lot of it was probably mixed up with intelligence derived from a lot of other sources. For example, we knew about the flying bomb and the V-2 before these things started to land on this country. It's worth remarking that the general philosophy of secrecy at RRE was that everyone on the establishment was entitled to know (almost) everything that went on. This was a particular philosophy of the chief superintendent, Mr. A. P. Rowe. The only exception to this was that there was a super-secret activity involved in countermeasures. The countermeasures group, which included people like Martin Ryle, was very secretive and it was difficult to find out even if you tried, which I didn't, to find out what they were doing.

Another question Dr. Aspray asks is what connection did the companies have with NRDC, National Research and Development Corporation.

There was, in fact, little or no connection between Lyons and the NRDC. The NRDC had been established as much as anything to exploit the patented inventions of Freddie Williams, (Professor F. C. Williams of Manchester University) and since we were not really following the Freddie Williams tradition there were only a very small number of circumstances in which we were involved with those patents. But there were, I believe, one or perhaps two patents in which we did infringe and I think we had to take out a license. I can remember, though not very clearly, some negotiations over the amount of the royalty that we paid. It might have had to do with Freddie Williams invention of the so-called B line which was the modifier count register, a feature which we incorporated into LEO I, having heard about it from friends in Manchester.

TAPE 1/SIDE 2

(I'm continuing to answer the question about NRDC): what is the assessment of NRDC's influence. I'm not in a very good position to assess the influence of NRDC on the development of the computer industry in Britain because they had very little impact on what we did at Lyons. I know they had a very considerable impact on what went on at Ferranti by guaranteeing the sale of some ten of Ferranti's first machines. The story is quite well known and I don't think I can add anything to it. As far as we were concerned, it was just a matter of agreeing the royalties that were due on those patents that they held which affected the design of LEO I, although they were important, these were not fundamental to the Leo design.

The tenth question in Dr. Aspray's letter is what were the business mistakes of Lyons when they sold out.

I think this is a loaded question. I don't think it was necessarily a mistake on Lyons' part to get out of the computer business. They could see that it was going to be a very voracious consumer of capital, since the rate of advance of technology would consume investment capital for a very long time to come; it might not show any return for many years. And, after all, they were not in that kind of business. So when they had an offer (which they hadn't solicited) from English Electric, it seemed a reasonable thing for them to do to agree to merge their interests -- especially since English Electric had an understanding with them that English Electric could buy out the Lyons interest completely at a later date, which indeed they actually did. So my assessment is that whatever funds Lyons put into the Leo project, they got back out again without significant loss, or, possibly, with a small profit when English Electric did take over the Lyons' interest in Leo Computers, Ltd. They did this about two years after Leo Computers was formed. My only sorrow is that as a director of Leo Computers, Ltd., at the time the merger was agreed I did not know that there was a secret clause agreed between Lyons and English Electric allowing English Electric to buy the Lyons interest out. But had I known I don't think I would have had any basis for objecting to it. The real question is whether it was a mistake for Lyons to have gone into the computer business in the first place; and I don't think it *was* a mistake. I think they made a unique and a very desirable contribution. Because after all, the people at Lyons knew better than anybody else alive at the time when the project started what was really wanted was to make a computer suitable for the office. They were able to demonstrate in the shape of LEO I, LEO II, and LEO III how those ideas should be put into practice. So I think they can claim to have established most of the standard practices in data processing by their own ideas and their own developments. Almost nothing of that came in from outside. So as far as the putting of ideas into practice is concerned, I don't think Lyons did make serious mistakes. They may have made a mistake in thinking that they would eventually be able to make money out of this but after all, if they got out of it without losing any money, which I believe they did, then that is surely some sort of justification for the present that they made to the rest of the world by developing the techniques of data processing from the beginning. However, I am aware that there is a recent paper by Dr. John Hendry of the London Business School, which he showed me before it was published; though I'm unfortunately unable to give you the precise reference to the paper in its published form [to be published by MIT Press]. He wrote a paper of some length analyzing the early history of Leo as a business enterprise and came to the conclusion that it made a great many mistakes from a business point of view. So if you tested the activity of Leo computers with the hindsight of a graduate of the Harvard Business School I have no doubt that you would find a great many mistakes. Perhaps the most fundamental one was that it was not looked upon as a money-making exercise in the first place. Of course had it been, I don't suppose they would ever have undertaken it or started on it in the beginning.

A further question is the association of Lyons with other companies through the Office Management Association (OMA).

I think this is really quite an important consideration. As far as I understand it, the story goes roughly like this. By about 1935, when I think the OMA was actually started, Lyons offices had got into a fairly highly organized state and in a sense there wasn't very much of the world left to conquer inside Lyons. So Mr. Simmons and his fellow managers looked around and saw that a great many companies were not managing their offices nearly as well as Lyons was managing its offices. So they got together managers from a great many other companies including such well known companies as W. D. and H. O. Wills, manufacturers of cigarettes and now part of Imperial Tobacco group; Stewarts and Lloyds who used to make steel tubes and are now part of British Steel, and about a dozen or more others; they started the Office Management Association. The idea was to exchange ideas about how offices were managed and in particular to carry out an annual survey of clerical salaries and grading; Lyons were very keen on grading office jobs and grading the workers who did those jobs so that they all knew whether they were getting a regular rate for the job or perhaps a rate at a superior level because they were doing the job particularly well. They all knew what the rates for the different jobs were. If you had a grading system which was uniform throughout the offices of a number of companies, it would be possible to see whether you were paying the rate for the job or much more or much less. So clearly it was a useful weapon in the office management armory. All of this was started by the Office Management Association; my understanding is that it continues without much of a break to this day. However, the Office Management Association has changed its name twice; it's now known as the Institute of Administrative Management, but it still flourishes. Mr. Simmons of Lyons was the first president. He remained president of the OMA until LEO Computers came into being, when he felt that as a person with a business interest in data processing he shouldn't remain as president because that represented a conflict of interest. But, of course, it was through the companies with which Lyons had become associated through the OMA that they found their first customers for the LEO's. So, for instance, a computer was sold to W.D. & H.O. Wills and another one to Stewarts & Lloyds, and so on.

When I first got to Lyons, the production of the annual clerical salaries analysis, as it was called, was a major event in the calendar; it was geared to the annual conference of the OMA. Though I'm sorry to say I never was invited to go to any of the conferences of the OMA. So I had very little contact, personally, with it. If you want to follow it up, they used to publish a magazine, I think it was called *The Manager*, and I suppose that some large libraries will have copies of this magazine on their shelves.

You asked me to try to reconstruct the whole Leo story, which certainly is very corrupt on pages 40 to 49 of the [earlier] transcript.

This is rather difficult to do if I try to follow exactly the pattern which was tried there, but I'll do my best. When I was first engaged by Lyons, almost the first thing we had to do was to see to what extent the EDSAC design would be suitable for clerical work and to what extent it would fall down. The most obvious respect in which it fell down was that it would not be possible to get numbers in and out of it nearly fast enough. Now it so happened that the previous summer, during 1948, David Wheeler, who was at that time a research student in Cambridge, was hired by Lyons for part of the vacation and he wrote a program in EDSAC code to do payroll. This was a program, that as far as I know, was never run; but it had at least one important consequence. That was that in writing the program, he and, I think, David Hemy estimated the volumes of data that would be required to be fed in and the volumes of results that would be required to be recorded for each man on a payroll and these numbers demonstrated that EDSAC was quite unable to cope with the flow of data that would be required. And don't forget that, knowing that EDSAC was not capable of handling great volumes of input and output data (after all it was a pretty slow machine by modern standards), they would not have included any more digits than were absolutely necessary. I have described the numbers in a paper which I think you have a reprint of ["Radio and Electronic Engineer" August 1975]. So there was this problem of the input and output and in the early months of 1949, when I was beginning at Lyons and before anything had started to be built, we cast around for a number of ways of speeding up the process of getting numbers in and out. There were really two problems. One was the purely mechanical problem of reading data fast enough from some sort of record and the second problem was how to convert the numbers into the binary form and to reconvert them after the calculations had been performed into a decimal sterling form so they could be printed. If you did this by subroutine, using the best subroutines that anybody then knew how to write, which early in 1949 obviously were fairly primitive, then as far as my recollection goes, it was going to take something like 90% of the whole working time of the computer to do conversion and reconversion from sterling or decimal form into binary and back from binary into decimal or sterling form. This was clearly unacceptable. We, therefore, came to the (probably incorrect) conclusion that we needed to have some external apparatus for doing the conversion/reconversion. Lyons as it so happened had not very long before this bought a telephone exchange from a firm called Standard Telephones and Cables. At that time, Standard Telephones and Cables were a subsidiary company of IT&T as was also the German company Standard Electric Lorenz and various other Standard Telephone or Standard Electric companies around the world. The telephone exchange that Lyons had bought was very satisfactory; Lyons were aware that STC had a research laboratory at Enfield and thought perhaps that they might be able to offer some help with this conversion/reconversion problem. As it happened at the time, STC were experimenting with a type of gas discharge trigger tube, which they hoped to use to design a telegraph repeater for repeating and regenerating telegraph signals sent from one teleprinter to another remote teleprinter. They were looking for some reasonably large-scale test bed for trying these tubes out; when Lyons came along they seized on the project very enthusiastically indeed. Now it so happened that they also had a very clever switching engineer, Mr. Esmond Wright, who was in the forefront of telephone switching engineers in the world at that time. He and Thompson got on extremely well. Intellectually, they were well matched. Both were extremely inventive. And in an amazingly short period of time, only a few weeks in fact, Wright and his colleagues had dreamt up the design -- the logic design -- for both converter and reconverter apparatus using gas trigger tubes. Our idea was that we would have buffered channels, buffers in several channels of input and several of output. The logic of LEO I lent itself to having up to four channels for input and four channels for output. We never actually fitted as many as that, but that was the maximum that could have been fitted.

The notion was that as numbers were converted they would be "assembled" in long delay lines so that you got a batch of sixteen separate numbers, one behind the other in a long delay line; you then gave a single instruction and read the whole batch into the computer very quickly in one operation. This meant that conversion and reconversion operations, as it were, overlapped on the operations of computing. In effect, the payroll calculations for one man were being performed while the output of the previous man was being reconverted and sent to a magnetic tape for recording and the input for the succeeding man was being assembled in the input buffer. Because there were several channels of input and several of output, one could separate the "current data," as Thompson called it, from the "brought forward data," which recorded what had happened the previous week, and from the "permanent data," i.e. data that did not alter or did not alter very often, such as a person's home address or whether or not he was a member of the pension scheme, contributory scheme for hospital expenses, and so on. Likewise, the output was split into two streams -- one to be brought forward the following week and one to be recorded and printed, say, on pay slips. So you needed, always, two streams of output and three streams of input. This represents the classical pattern of practically every batch data processing application that's been written ever since.

By about the end of 1949, therefore, the plans for STC to build not only the converter and reconverter equipments but also some associated and rather advanced magnetic tape recording gadgets were laid. Well, I'd better say something about the magnetic recording idea. They had conceived the idea of recording teleprinter characters which, as everybody knows, consist of seven elements (a start element, five information elements, and a stop element) serially on narrow magnetic tape. The tape was to be dumped into a box in such a way that it was never wound up, you just pushed it into the box. Because it was an endless loop, it could not get tied into a knot. As Maurice Wilkes once said, "You cannot tie a knot in two dimensions." So that, provided the layers of tape never passed one another, they could never get tangled, and curiously enough they didn't. A number of tape recording machines were made with tape pushed into a box like this but I think the STC ideas were as early as anyone's. I rather think that a similar type of tape recording machine was made by the National Bureau of Standards in Washington at about the same time or a little later, and certainly another design using a wider tape -- half inch tape -- was investigated and marketed on a small scale in this country by DECCA and ICT used it. Well, the STC tape decks did not work very well. In fact, they worked pretty badly. This was not because there was anything wrong with them in principle, but because they were not engineered to a satisfactory standard. I think had we not had trouble with the gas trigger tubes, then probably enough effort could have been put into the magnetic tape drives to make them work, though the arrangement was a fairly clumsy one. The original idea was that the input data would be recorded at teleprinter speed, which was the maximum rate at which you were allowed to type: namely seven characters a second. It would then be fed into the computer some six and two-thirds time faster than that; in other words, at about 45 characters per second. It would be converted and loaded into the long delay tubes so that you would get a long delay tube full of information in a second or two, and then the output would be reconverted by the reconversion apparatus as instructed either into decimal form or into sterling form. That would be recorded at a rather higher speed, twenty times teleprinter speed, in other words at about 140 characters a second, on one of two tape transports connected to each of the two output channels. When one box was full of data, it would be sent to the other machine to allow the box to be changed by an operator.

Now as I said, this equipment never actually worked properly. It worked for long enough to demonstrate the principle but there was a fundamental physical flaw in the behavior of the gas discharge tubes, not fully appreciated by Standard Telephones at the time, and which they sought to deny ever caused trouble. Unfortunately, evidence of trouble was only too visible to the naked eye so that you could not repudiate it. What happened was that a tube was supposed to be extinguished by lowering its plate potential sufficiently far and then when raised again, provided the trigger electrode was at a low potential, the tube should remain out. Very occasionally, perhaps one time in a hundred thousand or so, the tube would strike again if the plate potential was raised reasonably quickly after the tube had gone out. This was obviously unsatisfactory because if you built a trigger circuit consisting of two gas discharge tubes cross-connected one plate to the other plate by means of a capacitor, then if both tubes came on together there was no means of triggering either of them off. You could not reset that kind of arrangement. In any case, a trigger or flip-flop which gave both outputs up at the same time was clearly nonsensical. Associated with the single trigger tubes, which were often employed in pairs, as I have just implied, STC also made a type of multi-cathode tube which had ten cathodes. The notion was that one of the cathodes should have a glow on it. They all had a common anode, while the glow was supposed to step from one cathode to the next every time a trigger pulse was applied to a trigger electrode. These were similar to the Decatrons marketed by Ericcson, but the Decatron was different in as much as it conducted a much smaller current and therefore drive a diode logic gate directly. The STC multi-cathode tubes were designed to be connected to diode logic using metal rectifiers as diodes. Unfortunately, the STC multicathode tubes too suffered from a defect; occasionally you got a glow on two electrodes at the same time; say, perhaps, on electrode number one and electrode number six; as you pulsed it the two glows marched around simultaneously instead of one. This was also visible to the naked eye; you did not need even an oscilloscope to see that something was going wrong.

There was considerable delay in the delivery of the converter and reconverter equipments from STC labs, which at the time were at Enfield. They subsequently moved to Harlow. Lyons were getting very impatient. When the equivalent was eventually delivered, it was reasonably quickly discovered that it didn't work correctly. Then we had to scratch our heads and see what we would do to meet the situation if it never did work which is what actually happened. So we decided was that we would abandon the idea of the rather advanced magnetic tape transports and rely on classical paper tape and punched card equipment for input, and classical tabulating machines for the output and printers. Now, when we designed the original racks for LEO Mark 1, I had decided that we should leave space in each rack for additional units to cater for modifications. Lyons were a bit opposed to this but they gave in, and we also left room on every chassis for about four additional valves in case of modifications. This turned out to be the lifesaving feature. The first stage was to deal with the conversion and reconversion. Subsequent to the initiation of the program of Standard Telephones and Cables germanium diodes had started to appear on the market. This was very significant because it was realized by Ernest Lenaerts that if you could generate the binary equivalents of the numbers 10, 100, 1,000, 10,000, and so on, you could readily adapt the automatic control circuits that Maurice Wilkes had invented in EDSAC 1 to do multiplication, to do a conversion operation. This could be done by adding the appropriate binary equivalents of the various successive elements of a decimal number. Supposing the decimal number was, we'll say 57 (a favorite number of mine), you would first have to add in binary one, binary two, and binary four. You would discover that the "seven" had these elements in it by testing the successive bits of a four bit combination representing the first decimal digit, the seven. Similarly when you came to examine the second group of four bits, which represented the five in 57, you had to add in the binary equivalent of 10, no binary equivalent for 20 but the binary equivalent of 40. So, the circuits were very similar to those which were used for testing the digits of a multiplier. It simply meant that you did not add in the multiplicand. You added in the binary equivalent numbers instead. As there were a lot of binary numbers to be created the total diode matrix required to hold all of the equivalent numbers needed, as far as I remember, about 1,000 diodes. Now that was a feasible number using germanium diodes but it would not have been a feasible number to incorporate if we had tried to do it by wiring up hot cathode diodes. Up to that point, all the diodes in EDSAC I and all the diodes in LEO I had been made using hot cathode diodes which take up a good deal of room and quite a lot of power. When the timing was worked out, we found that we could do the whole of the conversion and the reconversion not only to and from decimal but also to and from sterling numbers, which were still in use and important at that time, in about 10% of the total working time of the machine. That got rid, at a stroke, of all the external converters and reconverters that STC had been trying to build but which didn't work.

The next thing to do was to find some way of feeding in the current data and the brought forward data. Here we used a combination of paper tape for all the fresh initial data, which was *recorded* using ordinary paper tape punching equipment and *checked* by some special equipment that we designed, which I will say more about in a minute. We bought tape readers from Ferranti, who by this time had put on the market quite an efficient paper tape reader which operated photoelectrically and read at 200 rows per second. It would stop and start between one row and the next. This was considered quite fast until later on another type of paper tape reader was introduced based on a design invented in the Mathematics Laboratory at Cambridge operating at 1,000 rows per second. However, 200 rows per second was quite satisfactory for us. But the "brought-forward" and "carry-forward" was handled using punched cards. The idea was that each row on a card would represent one binary number of 35 bits. Thus one punched card would carry 12 binary numbers. It works out that even though we only used fewer than half the columns on the card (which, of course, carries 80 columns), we could put onto it in binary form the equivalent of about two and a half times as much data as it would carry in conventional decimal form. That meant that cards were reasonably efficient as carriers of information although of course the arrangement was completely and absolutely non-standard. Nobody else used cards in this way at the time and nobody else used them in that way afterwards, as far as I'm aware. But it meant that by coupling the card readers and card punches to the buffering arrangements, which we originally intended to use with the Standard Telephones converters and reconverters, we could easily feed data in and record results for carry forward purposes on punched card devices without holding up computation.

I might say that we were quite conscious of the dynamics of programs which was made audibly apparent by my idea of connecting a loudspeaker to the machine. When the computer was carrying out instructions in the program, it made a very erratic scratching sort of noise because each instruction took a different time from the preceding and succeeding one. However, should one of the input or output channels get blocked because, for instance, one had halted the card reader to allow more cards to be loaded, then the computer "marked time," which produced a steady continuous note. So you got this scratching noise interrupted by short or long continuous notes which sounded very different and gave an extremely effective audible indication of the extent to which the computer was being held up by input or output transfers that didn't overlap regular computation. Programmers found this a very easy way of judging whether they had balanced their programs properly or not. All these developments were put in hand and were made effective while STC was still struggling to get their equipment to work. But by the summer of 1953, it was clear to Lyons' management that they never were going to get it to work and so, very sadly, Lyons told STC that the business arrangement was at an end and they could take the equipment away. They agreed to pay what had been originally agreed, although in fact it turned out that STC had expended probably about five times as much and pleaded with Lyons to be reimbursed to a greater extent. Lyons, however, were adamant and STC had to put up with the original sum promised, which was about twenty thousand pounds if I recollect correctly. So by the late summer of 1953, we were really ready to start planning to put work onto the computer.

TAPE 2/SIDE 1

Lyons' attitude was that if you were not doing work that the business depended on, then you weren't really using the computer operationally. It was true that since 1951 a small job had been run regularly. This was a bakery sales analysis; originally it ran extremely slowly when we had a single teleprinter for output. Eventually we could do the whole job in about fifteen minutes when we had the tabulators connected to provide a printed output. At the beginning, it took a great deal longer than that, sometimes several days, because of computer breakdowns.

At this stage, I think I would like to say something about the maintenance arrangements on LEO 1 which, I believe, were considerably more sophisticated than on most computers at the time. We had adopted the principle of "marginal testing," which we heard about through Maurice Wilkes, who I think had also probably heard about it from somebody else; we applied it in a very systematic way. The notion of marginal testing is that you try to anticipate the effects of the gradual decline in valve emission and valve mutual conductance by deliberately pushing the circuit a little bit off balance deliberately to see how far you can go before a test program demonstrates that a fault has occurred. If you can push it so far, then you may reckon to have a working margin and you can argue that the machine will be good for at least another 24 hours use before something goes wrong on an operational job. So the idea was to apply a suitably graded 50 cycle[*] signal to vary the bias in the grid circuits of all valves connected as discriminating amplifiers.

Most of the amplifiers, indeed all the amplifiers I think, in pulse circuits in LEO 1 were long-tailed pairs, the principle of this type of circuit is that if there is a sufficiently small signal, nothing comes out and if there is a sufficiently large signal, a saturation signal comes out. If you increase the bias, the input signal required to produce saturation signal will be bigger. So you have, in fact, *reduced* the margin by *increasing* the bias. By using an alternating fifty cycle signal you both increase and reduce the bias so increasing the chance of a breakthrough signal, which should have been wiped off, emerging and so affecting some subsequent circuit, and also increasing the chance that a wanted pulse signal which is on the edge of being too small will actually fail to operate a succeeding circuit. So you test both margins at once by using a fifty cycle signal.

We had an arrangement on LEO I whereby the magnitude of the fifty cycle signal applied to practically every one of the places in the machine (and there were several hundred of them) whose margin was to be tested could be adjusted individually to suit that circuit, which meant that we could apply the same marginal test voltage through the entire machine and yet test each circuit to a graded extent. If any circuit in the machine had declined since the last time we did this, perhaps 24 hours before, then we knew that that circuit was more marginal than it had been and perhaps a

[*] Derived from the public supply which in Britain operates at 50 Hertz -60 Hertz in North America.

valve needed to be changed. That was the principle of marginal testing which we found to be very effective. Indeed, we copied it in LEO II, and it was equally effective in all the LEO II's. It is not so effective a scheme when you come to deal with transistor circuits because transistors don't suffer from the same progressive loss of emission as do valves; the types of failure that you get in transistor circuits tend to be of a rather different character. So it has tended to go out of fashion. But we did have a very considerable armory of test programs for Leo 1 and on the whole we reckoned we could find a fault, particularly a catastrophic fault, within an hour or so. Indeed, I think the machine was seldom out of action for as much as 24 hours at a stretch. It would have been considered shameful if we had not been able to mend it considerably more quickly than that.

Another form of marginal test which we applied as a routine to LEO I and to LEO II was the store frequency margin. In the original EDSAC the store frequency was adjusted by hand. In the NPL ACE machines, which also used mercury delay lines in a thermostatically controlled enclosure, a crystal controlled oscillator set the repetition rate, and was capable of a very small range of adjustment. Once it was adjusted, provided the temperature of the delay lines didn't change, there was no further need to alter it. In LEO I and II, however, we employed a different system. We decided that we would use one delay tube as a reference line and use that to control the appropriate frequency by using a reactor valve across the oscillator producing the clock pulses. So as the temperature of the delay lines varied, so the clock pulse could be varied in sympathy by a d.c. control voltage. A control in the feedback loop that produced the control signal also allowed you to adjust the relative phase of the clock pulses so they could either be a little early or a little late; that could give you a measure of the so-called "range" of the store. If all the delay tubes were not at exactly the same temperature, then some of them would lose pulses during a store test more readily than others; thus you could calibrate a knob to see how much variation in the frequency you could stand before any of the delay tubes started losing pulses as evidenced by the store test. This was regularly used for testing the store both on LEO I and LEO II. It was found to be very effective.

Another aspect of the maintenance of the LEO 1 was that every thermionic tube or valve was identified by a number painted on it; its life history was recorded on a card so that we knew exactly when it was put into service, where it was put, and when it was taken out of service. The aim was to have every valve tested for its emission and other characteristics such as interelectrode insulation and so on at regular intervals. I can't now remember exactly how often we did test the valves but my impression is that we endeavored to test each of them every few thousand hours. On the whole, we did not expect a valve to last much more than ten or twelve thousand hours. We took expert advice on valve life testing from Mr. Brewer of the British General Electric Company Research Laboratories at Wembley, who was in charge of the valve life test laboratories. By curious coincidence he happened to be a next door neighbor of my boss, Mr. T. R. Thompson, which was how we came to get in touch with him in the first place; we found him extremely helpful.

This is the stage in which to say something about the numbers of people involved in the project. Up to the end of 1953, there was only one computer and, therefore, only a very small number of people were involved. My impression is that the total would have been around a couple of dozen about, equally split between programming and engineering staff. Once the LEO I went into service regularly in the spring of 1954, it did a payroll for Lyons in January that year and it never failed to produce it on time during the week. By the summer of that year, it was also doing a substantial part of the Ford Motor payroll for Ford Motors' works at Dagenham, and that also never failed to produce the payroll for them during the course of the week. Which, considering the average time between failures too, the machine was probably somewhere out around 10 hours on average, was, I feel, an incredible achievement.

At that stage, Lyons were looking ahead; they started thinking in terms of a second machine, which we were going to call LEO II. Now here, I think, we probably made a few engineering errors. We were still thinking in terms of a valve machine because transistors were still not commonplace. They had been invented by then but nobody would have thought of building a computer with them in 1954. We did, however, want a machine to run faster, and it occurred to me that one way to make it faster was to speed up the store without altering the speed of any of the rest of the system. It seemed to me that it would be possible to exploit the intrinsically wider band width of a quartz mercury delay tube by using a higher pulse repetition rate in the store than was used in the rest of the machine. So we devised a store using an inter-learning arrangement. The pulse repetition rate was actually phase-locked to the radio frequency signal, which it had not been in LEO 1. So we started off with an oscillator at about half a megahertz and multiplied its frequency so as to produce a carrier signal at 13.5 megahertz. That signal was switched on and off into the balanced transmitter that was used for sending pulses into the delay line. If you used quarter-microsecond pulses, as we did, you were going to need a band width of something like 6 megahertz. We therefore decided to use a single sideband arrangement rather than put the amplifier pass band symmetrically above the radio frequency band width. A tolerance calculation made clear that we would need to control the lengths of the delay tubes, the band width of the amplifiers, and the lengths of the connecting cables all to within fairly tight tolerances; and this we actually did. But it turned out to be a good deal more difficult than we had expected to design the so-called storage unit that recirculated quarter microsecond pulses in place of the one microsecond pulses used in LEO I. I think the moral of this is that you should never push any technology to its limit, but stop about half way as far as it will go; by the time you get it all to work someone will have invented a technology which is intrinsically faster anyway. This kind of argument was reinforced by the experience of English Electric with the KDF9 machine; its sophisticated logic made a 6 microsecond store look and feel like a 2 microsecond store. Again, by the time it had been made to work you could buy core stores with a 2 microsecond cycle time.

LEO II's were designed to be manufactured and sold to other people and in the event eleven of them were manufactured, no two exactly the same. Their fundamental concepts were much the same as LEO I except for the fact that because of the interleaved pulse arrangement in the store they ran appreciably faster; I suppose about two and a half to three times as fast. The arithmetic circuits were no faster, but the store was because the delay lines were only one quarter the length. A number of other physical changes were made to the machine. The units were mounted vertically instead of horizontally, and we rearranged the power supply distribution arrangements and the fusing and a great many other things to avoid some of the electrical risks we found we had experienced with LEO I. But the most important change was that, after the first two or three, the machines were fitted with tape transports which we bought, initially, from Decca, who made a twin tape transport operating on half inch tape, which had the sort of performance of modern tape decks available from other companies around that time. Some of these machines were also fitted with magnetic drums, which we bought from Ferranti. The logic of the connection to the Ferranti drum was rather sophisticated because the drum had to revolve at a specific speed and transfer information to a synchronous machine; however, we finally made this work quite well. Several of the machines in the LEO II series were installed with Ferranti drums.

The biggest problem, from an engineering point of view, we had was the famous Samastronic printer. Powers-Samas, which had been a punched card company trading for years on the design of mechanical punched card equipment that did not use electrical sensing of the cards and did not use electrical methods for accumulating totals read from cards, decided that they would make a quantum leap and to overtake the engineering practice of the British Tabulating Machine Company, which did use an electrical sensing of cards following the IBM pattern. Powers introduced what they called the Samastronic series of punched card equipments. The "tronic" bit was supposed to indicate to people that it was electronic, which in fact, it wasn't. It was just as electrical or electromechanical as the despised British Tabulating Machines products. However, the irony of it was whereas the mechanical engineering of the punched card equipment from British Tab was actually superb, the electromechanical engineering of the new Samastronic equipment from Powers was rather bad; it was really a very unreliable piece of machinery. I don't think the tabulators really ever worked properly at all. The printing unit was extremely ingenious in its concept, but badly executed. The idea was to use single stylus wires vibrated backwards and forwards across each printing position, in such a way that had it been in continuous contact with the paper it would have produced a kind of sin wave, in fact a black bar right down the paper. The paper was moved steadily upwards at the same time as the stylus was wobbled backwards and forwards from side to side. By impulsing the stylus at the appropriate moment on its sinusoidal track you could put dots onto the paper and so build up a picture of the character in the form of a matrix in much the same way as modern matrix printers do using a large number of separate styluses. However, this printer used one stylus only for each print position. As it had some hundred odd print positions across the paper, there were, therefore, approximately a hundred wires actuated by small moving coil loudspeaker elements, at the ends of the wires. There were many reasons why this machine was mechanically unsound, which it would just be boring to recount; we found a way of driving it electronically rather by brute force using a fairly large output tetrode with 800 volts on its plate. That gave enough power to drive one or two coils in parallel. They were intended to be driven from a 20 volt generator through an enormous electrical commutator; the electronic system, as far as that went, was quite reliable. The printers printed at 300 lines per minute, which at the time was considered extremely fast. In practice if you had two of them and a team of skilled mechanics, you could have at least one of the printers in operation for most of the time, which indicated the unreliability that they suffered from. In spite of all these problems, the Leo II's gave their purchasers a reasonable degree of satisfaction and, as I said, approximately 11 altogether were made of which 10 were sold as Lyons kept one for themselves.

Towards the end of the series, we had learned enough about the transistor circuits and about the design of core stores to embark on the design of a core store of our own, which we substituted for the delay line store. So there was a model called LEO IIC, which had a core store of approximately 8,000 words, which was eight times the size we could afford to provide with delay tubes. This was our introduction to transistor circuitry and also to parallel circuitry because, of course, the core store was parallel. However, there were still delay lines in the LEO IIC machine for the individual computing registers and for the input and output buffers. So LEO IIC was a sort of hybrid machine, using mainly valves with quite a number of individual transistors in the core store.

I think in passing I really ought to pay tribute to Stuart (?) Williams, who was working for Telemeter Magnetics at the time. He came to see us at the factory that Lyons had established at Acton, and in the course of about two hours he told us what we needed to know about the design of core stores allowing us to do the engineering ourselves. I don't think, perhaps that this was his intention, but it was in fact what happened.

During the second half of the 1950s, we were working at a factory on the North Acton Industrial Estate, which was set up to manufacture small numbers of LEO IIs. The publicity film made in 1957, gives some idea of the atmosphere in which we were working and the general conditions in the factory. Towards the end of that decade, it became quite obvious that you could design a parallel machine using transistors, which would be altogether faster, more reliable, and generally superior in every way to any kind of valve machine. So, we set about planning to build such a machine, to be known as LEO III. Originally, the core store was to have a cycle time of around about 12 or 13 microseconds, but later versions were built with faster core stores. It was going to have a microprogram for controlling the way the arithmetic unit worked following Maurice Wilkes' ideas at Cambridge. It was going to have multiple, buffered input and output channels just as in LEO I and LEO II, but, in this case, we realized (without having to be told about it by IBM or anybody else) that we could incorporate the buffers in the main store. Therefore they could be as large as anyone wanted within reason and there could be as many of them as anyone wanted; in practice

I think two were considered by the programmers to be sufficient for most purposes -- two on each channel that is to say. But they were certainly able to be as big as you wanted. We conceived the idea of autonomous controllers for the channels, which would build up blocks of data in store as it was read, from a magnetic tape, a punched card a paper tape, or whatever; likewise data was "taken to pieces" for dispatch to the printer, the magnetic tape recorder, or whatever output device was involved.

The one thing we did not recognize was that the interface between these channel controllers and the main machine should not have been placed where we actually put it. We located the interface on what you might call the main store bus so that the device controllers had to calculate the store addresses, and the whole interface used a very large number of wires -- I think some 60 odd wires. What we should have realized was that a character-wide interface, without requiring the store address to go to the individual device controller, would have been quite sufficient to carry the data traffic; that would have separated the channel from the device control in the way that IBM did in the 360 series. This was a trick we did not invent, unfortunately. We came to the knowledge of what IBM had done in the 360 series after we'd taken all the design decisions for LEO III. Of course, this is what always happens!

The design of LEO III was proved to be correct by building a small-scale prototype of 13 bits only. The whole machine used 40-bit words, plus additional bits for various other purposes. For example, we had a store reservation scheme that, as we discovered afterwards, was very similar to the one IBM invented, except IBM reserved the store in blocks of 1,000 words at a time, whereas we had reservation bits associated with every separate word in the store. This was rather extravagant and made the stored words very long; including parity bits it was 48 bits long, which I suppose was really rather extravagant.

The microprogram which controlled all the arithmetic operations and also did the fetching of the next instructions and so on, was wired onto a very sparsely populated core matrix. The idea was to place core wherever you wanted a "one" output and not to have any core where you didn't want one. Some hundred odd readout wires of them controlled the various gates between registers through which signals were to pass. My recollection is that there were 32 stages in each so-called core plane. If we had a rather complicated microprogram, for say doing a floating point operation, then more than one core plane had to be employed. Unfortunately, there was no full description of the design of the microprogramming arrangements in LEO III in print. This is a pity because I actually think that LEO III was the first machine to go into production, using Maurice Wilkes' ideas, although we gave effect to them in a way very different from that in EDSAC II.

As a result of the experience with LEO III, we became extremely keen supporters of microprogramming as a technique of designing computers. It allowed you to overlap several stages in the design process and change your mind without incurring disastrous consequences if there was something wrong with the detail of the logic design of a particular instruction, say, or if the programmers asked you to put another instruction in or to change an instruction that was already there so that it was operated in a different way. I would claim that LEO III was an extremely successful design. Of course, it became rapidly outdated when integrated circuits came along since it used individual transistors. By comparison with later designs it was extravagant in the number of cabinets required and in the number of boards required to carry the logic. However, it was a very big advance on LEO II, and it was very much more reliable. In consequence, its not surprising that we managed to make and sell about 150 of them, which gave very good service in this country for quite a long time. Several were exported to other parts of the world, notably to South Africa and Australia.

One of the more interesting things about the LEO III was the opportunity that it provided for multiprogramming or sharing the computer's time between several programs at the same moment. This was an idea that we stole, unashamedly, from Ferranti. But I believe we actually made it work before Ferranti did because they were having some trouble with circuits at the time, and we didn't have any circuit troubles only programming ones. One of the most notable examples of sharing the computer between several programs was the suite of seven programs that were developed by the then Post Office for telephone billing in the London area. This was part of what is now BT. The seven programs were very tightly interlocked in such a way as to use up all of the available computing and input and output facilities of LEO III. The suite (?) ran every day of the year for subscribers for telephone service in the London area. Two LEO IIIs at least were devoted to this job and did nothing else. When the LEO IIIs went out of service and were later on replaced by System IV machines made by English Electric Computers, the Post Office (?)

begged us to arrange for the System IV machines to execute LEO III codes so they wouldn't have to reprogram this suite. This was done. Even more remarkable was that when they subsequently bought 2900 series machines, which were also microprogrammed, they once again asked for the 2900s to be microprogrammed to execute Leo III code so that they could continue to use this same suite of programs. And as far as I know, they may well be in use to this day. This seems to be an example of John Brown's body, moldering in the ground while the soul goes marching on.

I want to talk about the LEO III magnetic tape system, which had several very interesting features, some of which have never been replicated on anybody else's magnetic tape system even to this day as far as I know.

TAPE 2/SIDE 2

I said I would record some details of the LEO III magnetic tape system. This used Ampex decks, but the sophistication that was built in had to do with the way the information was checked and was subsequently recovered from the tape, which we thought was rather superior. It was based on what we called "two-level discrimination". I believe that a somewhat similar scheme was devised by IBM, though I'm not now quite sure how they applied it. Our idea was to ensure that recorded tapes were (a) checked for connectors and (b) also checked to see that they had sufficient margin built into the recorded information to guarantee it could be recovered later. This was done as follows: when signals were read from the tape, they were fed in parallel to two different discriminator circuits -- one with a level set rather high and the other with a level set rather low. The high-level discriminator would not respond to signals that were too weak and the low level discriminator might respond to breakthrough or noise signals if they were too strong. The idea was that signals passing through both discriminators should be checked and that on recording, the levels were set further apart than they were on reading. On recording, a parity check had to be effective on every row of the tape, and also there had to be a correct version of a sum check for every block on the tape. This meant that the entire block was in effect checked separately through both the high level discriminator and the low lever discriminator, which were themselves further apart than they would be when the tape was read. When we came to read any tape, which had been successfully recorded in this way, the two discriminating levels were rather closer together. I can't remember, without reference to the files, exactly what the levels were; but suppose that on recording, the levels were, say, 30 and 70 percent, then on reading they might have been 40 and 60 percent. On reading, however, we took the signal from the high-level discriminator, granted that parity for that row was satisfied; if it were not we took the version fed from the low level discriminator, granted that that was satisfied, and we then performed the sum check at the end of the block. We thought we had a greatly increased chance of reading successfully a tape that had been passed through the checking process in the first place. If on check reading a tape -- every block was check read of course -- the block was not successfully recovered then we merely canceled it and recorded it further down the tape. My claim is that with this system we never had a block read back that was passed by the checking circuits but which was subsequently demonstrated to be faulty by some program check. As far as we know, this never occurred. It certainly did seem to enhance the readability of tapes, which at that time were not of anywhere near as good a quality as the tapes manufactured today. They probably had more bad patches on them or dropouts on them. I consider that that was a very satisfactory system of reading and recording tapes.

I'm now going to try and answer Dr. Aspray's question 15, which has to do with my position in product planning at English Electric.

At the time of the merger with English Electric, the design of LEO III was already settled and it was in production. We were thinking in terms of another machine and we agreed with English Electric that a machine should be planned to be a replacement for both KDF 9 and for LEO III. The managing director of English Electric computers, from the point at which it was formed, was Mr. Wilf Scott, he thought that we should be able to produce a product plan by Christmas 1963. I was given about two staff to help me do this; we failed utterly to produce anything concrete in that time. As a result, the research activities and the product planning activities were separated at that point, and product planning was moved to Kidsgrove from Minerva Road, Acton, where I had been trying to run it. A larger team of people got to work on it. However, even they were unable to complete the product planning for a new range sufficiently quickly. Mr. Scott, naturally enough, became very impatient. So by the end of 1964, during which year, of course, IBM announced their 360 range, it became increasingly clear that some crash decision would have to be made. A secret conclave was called over the New Year in 1964/65 in which some of us were shut up, like Cardinals to choose a new Pope, in English Electric House and told to make up our minds on the basis for the future series for

English Electric. We were really faced with three possibilities: to base it on the LEO III design which was working successfully, to base it on KDF 9 which was also working successfully, or to take advantage of an information sharing arrangement which Marconi had had for many years with RCA and base it on the Spectra 70 series. This had been announced earlier in 1964 but after the announcement of the IBM 360 series, earlier in that year. A consideration of the effort that would have to go into software and the prospect of sharing this effort with RCA, caused us to decide to support a scheme based on the Spectra 70 series, to be known as System IV. As everybody knows, the Spectra 70 series was compatible as far as the ordinary instructions went with the IBM 360 code but was not compatible in respect of the privileged instructions. Although the privileged instructions were different, RCA devised improvements over 360, which meant that although Spectra 70 couldn't have the same operating system as the 360 series, it had many advantages over 360. There were four machine states instead of one, which meant that when you had interrupts to enter the operating system the details had to be put away in store and restored after you had dealt with the system call. Ironically, the argument that we would save on programming efforts and costs by sharing them with RCA was not borne out in the event because we got practically no programming benefit from sharing the work with RCA. Had we done it entirely on our own account, we would probably have had to put in just as much effort and we'd have probably done it just as quickly in the end. You will find some part of this story told in an article in the May issue of the *ICL Technical Journal*, May 1988, written by Martin Campbell Kelly who got a lot of the information from ex-Leo and ex-English Electric people, of course.

One of the questions Dr. Aspray asks is: Can I describe the competition in relationship of English Electric to IBM and RCA?

Well, I think to answer that question properly means going back to the origins of RCA, and a lot of what I say now may be hearsay so may need checking from other sources. But I believe it to be essentially correct. At the period just after the first world war, the British Marconi Company was also active in the United States and RCA did not exist. The American government was concerned that such an important aspect of military operations as radio communications should be largely in the hands of a foreign-owned company; namely, the Marconi Company. So they insisted on Marconi selling its American interests to an American company, which was then called the Radio

Corporation of America. The result of that deal was that an information-sharing agreement was entered into between the British Marconi Company and the RCA Company, which have persisted for many years. Before the war (according to my information), RCA had made a great deal of its money from licensing other people to use its patents and its designs. In particular, the famous American range of thermionic valves (as we used to know them) were in fact RCA designs, manufactured to RCA specifications by a great many other companies in the United States under license and sold around the world. We had no comparable arrangement in the UK, so American valves tended to dominate world markets; all were licensed by RCA. I think RCA got quite a lot of money from this. So RCA had got into the position of licensing other companies widely and extensively. But Marconi preserved an especially favored license agreement, which was inherited by English Electric when English Electric bought out the Marconi Company in the post second war period. This made it possible for English Electric to acquire information about the Spectra 70 series under this information-sharing agreement on very reasonable terms and for a very much lower cost than ICT was paying for very similar information about very similar if not identical machines.

I can't think of any meaningful way of discussing competition between IBM and English Electric. There wasn't really any sensible competition between the two except in the UK; as far as the American computer market was concerned in the 1960s, when RCA were active in selling computers, obviously IBM had a very much larger share of the market than RCA ever had. This, I suppose, was really due to the fact that IBM had dominated the punched card market in the first place and if people were going to move from punched card machinery to computers, it was quite natural that they should stay with IBM; provided IBM were able to satisfy them with efficient and up-to-date computers, which of course they did. What really ought to be a matter for surprise was that, once IBM had decided to go into computers, anybody else was able to sell computers in competition with IBM at all,

The last question in Dr. Aspray's letter is to describe the end of my career with English Electric.

In many ways, this was rather a disappointment. I got detached from the product planning activity once it had been decided to go ahead and base the new range on the Spectra 70 series. The detailed design of the machines was done either at Kidsgrove or there were two smaller machines, which were not successful, which were designed at Marconi's at Chelmssford. This was, to some extent, due to Sir Gordon Radley, who was a Director of Marconi and anxious that Marconi should have a hand in the computing activities. But the machines that they were invited to develop would not have been successful in the market place because they really hadn't got sufficient power. This, I think, was a planning mistake. In order to be able to sell them more cheaply they were degraded machines with insufficient of the Spectra 70 facilities in them to make them effective. I concerned myself with research and got very interested in the use of quite advanced physical techniques to create and interconnect circuits. I was convinced that interconnection was the key to machine design, which I'm still sure it is, but I didn't realize quite sufficiently how much of the interconnection would be done on a chip. We were rather obsessed with the problems of interconnecting and coding chips. I came to the conclusion that one of the ways this should be done was by means of an electron beam machine. Somebody, I'm not now quite sure who, dreamt up the idea that if you used an electron beam machine, you could melt a small blob of, say, solder through the chip by turning it upside down and do the fastening of the chip to a substrate by having an electron beam pass right through the chip at the edge where it wouldn't do any harm. The heat dissipated in a solder blob would fuse it in approximately a microsecond. So in a very short space of time you could stitch down a chip onto a rigid substrate preprinted with the appropriate wiring. In any case, the argument was that an electron beam machine was a very powerful tool for a great number of purposes and for a great number of reasons. One of which was that not only could you apply power and mark things, heat things, and etch things, but you could actually use it as a scanning microscope at the same time to examine what you were doing. On the strength of all this, we built elaborate laboratory facilities at Minerva Road, but these were overtaken by the merger; ICT did not consider that we were doing anything of great value. In fact, ICT were to some extent antipathetic to the investment of effort in that sort of advanced technology, so the whole of this laboratory equipment and the team was dismantled and disbanded. This was a big disappointment because we'd actually got an electron beam machine on order and partly manufactured by a specialist firm in West Germany. I'm reasonably convinced today that it would have provided valuable experience. In fact, I see that the firm European Silicon Structures are using electron beam methods for manufacturing special customized integrated circuit chips. So I'm still a believer in the use of the electron beam machine. But all of these ideas were ahead of their time and the management of ICT were not particularly in favor of funding research on a generous scale, since they foresaw having to write off large losses that English Electric computers were making around about this time as a result of the merger which was wished on them, in some ways, by Mr. Tony Bonn -- the Minister for Technology.

After the merger took place, I fairly quickly (within about 18 months) gave up any responsibility for research at Minerva Road and moved into the team at Putney planning the 2900 series. I had a number of different responsibilities. One of which was to try to express thoroughly the requirements for maintaining the machine. There was quite a large team of people doing this planning operation -- some 60 people -- and so the jobs were parcelled out.

I would like to speak, for a little while, about my interest in standards because this has long been an ongoing interest. Even after I have retired, I have had an interest in activities in the standards field. As far as the computer industry is concerned I suppose, it started in about 1960 when there was a move to form the European Computer Manufacturers Association. The idea was promoted by three companies; namely, ICT, IBM and Bull, whose respective European chief executives wrote to all the other computer manufacturers in Europe -- there were about 16 of them at the time -- and said wouldn't it be a good idea to form a standards-making organization to take account of European requirements. Otherwise, the standard scene was tending to be dominated by American ideas; and Americans didn't necessarily know the specific requirements of European computer manufacturers. The other manufacturers all thought this was a good idea and a meeting took place in Brussels in 1960 to which I was sent by Leo Computers. Agreement was reached in principle to form the Association. Further meetings ensued to all of which I went as the LEO Computers representative and I became the official Leo Computers representative to ECMA, as it was then called. I continued to be the English Electric representative after the merger with English Electric, and I continued to be the ICL representative on the formation of ICL. I represented ICL at ECMA until retiring in 1984. Thus, I had some 24 years of activity in the European standards field, and I am very happy to record that. I am also happy to record that the members of ECMA elected me president on two occasions, for two years on each occasion. I also served as vice president and as treasurer. Since I retired, I was asked by BSI with government support and together with another consultant, to produce a review of the policy that should be followed by the British Standards Institution, which at that time, about three years ago, was faced with very much more work than it could properly support through its limited secretarial staff. We did produce a report, which has had some beneficial effect, though it hasn't had all the effects that we were hoping for.

After the planning of the 2900 series was completed and it moved into the design phase, I tended to move into a marketing role, which was a somewhat surprising position to find myself in. But, looking back on it, my main contribution was that I was able to explain to marketing managers the technical implications of many ideas, which were otherwise not totally obvious. For instance, I produced a report, which (I like to think) was responsible for ICL changing its policy on office automation; and I also produced a lot of material and reports on how ICL should tackle the problem of marketing in countries where the language spoken was not English and where the customs and practices were not those which obtain in the United Kingdom. I think these activities have had, over the years, a significant effect on the way ICL has done its business.

For a period I was also very interested in establishing uniform measures of performance inside ICL so that one computer could be compared with another and so that we could assess the performance of our products in comparison with those of competitors. I was very considerably assisted in doing this by a very clever colleague, who is an extremely good mathematician, Conway Berners Lee, who of course like myself is now retired. But we did establish a uniform approach to performance measurements and performance comparisons in ICL in the middle 1970s, and had a significant effect on the way marketing plans were made.

I will wind up by recording very briefly what I am still doing. I suppose more than half my time is devoted now to acting as a series editor for a publisher of books on information technology. This is a small British company called Ellis Horwood Ltd. of Chichester in West Sussex, who are closely associated with John Wiley. John Wiley distributes their books throughout the world, but Ellis Horwood take the publishing risks in the first place. They have a series on information technology of which I am the main series editor. I receive manuscripts and proposals for books at frequent intervals, and I'm asked to advise what should be done about these books: whether to publish them or not and if they are published whether they are satisfactory in the way in which they have been prepared by their authors. I have found this an extremely interesting exercise and it couples with the efforts I put in for ICL sitting on the board of their Technical Journal, with which I have been involved since it started about eleven years ago.

I also act as a monitoring officer for the British Department of Trade and Industry, which supports what used to be called the Alvey Research Program. This is a large-scale computer research program initiated in competition with the European Esprit program six or seven years ago. It's named after the Post Office engineer, John Alvey, who chaired the committee that recommended the program. In recent years the government has tended to put more money or to encourage firms to put their efforts through Esprit rather than the Alvey program. But this program still continues, and the particular project I am interested in has to do with the human/computer interaction and the adaptation of computer behavior to the behavior of people. This is a very interesting subject, although I can't claim any particular specialist knowledge of it.

END OF INTERVIEW