A COMPUTER ENGINEER LOOKS BACK

Maurice Vincent Wilkes

Thank you very much.

My earliest distinct recollection is of a small girl running up the garden path of the house where I lived with my parents, eager to share good news. The good news was, “There's peace.” It was the 11th of November, 1918. The First World War was over. I was five years and four months old. When I was six I started school. A slow start perhaps, but I soon began to find school work interesting and congenial, and I cannot remember a time when I did not want to be an engineer or a scientist of some sort. When I was twelve it was a great joy to me to start the study of physics and chemistry along with mathematics. By the time I was fifteen I started calculus.

In 1931 I went up to Cambridge University, where in my undergraduate course I had a good grounding in mathematical physics; notably mechanics, electromagnetism, thermodynamics, and similar subjects. Also wave mechanics, which at that time was only a few years old. My education in pure mathematics would, I suppose, now be considered a meagre one. Mostly analysis, geometry, no more than bare elements of group theory, and no set theory.

I took my bachelor's degree in 1934 and joined the Cavendish Laboratory as a research student. Lord Rutherford, the great Lord Rutherford, was in the middle of his tenure as head of the laboratory, and atomic physics was going through one of its greatest periods. I did not see too much of that work because my research was on the propagation of radio waves in the ionosphere, one of the few nonatomic branches of physics that Rutherford supported in his laboratory. I followed, to some extent, the work that was going on in atomic physics but I have since regretted that I was not sufficiently close to the work of my fellow students to share in their excitements.

I’d been attracted to radio-wave propagation by a long-standing interest in radio. I'd held an amateur transmitting license for some years. I'd designed and built a good deal of receiving and transmitting equipment. The mission of the Cambridge radio group was to study the characteristics of the radio waves that were reflected from the ionosphere, in the hope that it might be possible to draw conclusions from them about the composition of the high atmosphere. That is, its chemical composition as well as its state of ionization. This involved the designing of radio equipment for measurement instead of communication—something that was rather new. The radio art had then only progressed to the point at which one could build equipment for measurements of that kind.

I took my Ph.D. degree in 1937, and it was around that time that I first became involved in computing. In 1937, there were no digital computers in the world. The term computer meant a human being working with a desk calculating machine. And people today are rather amused if you use the word computer to mean a person and not a machine. However, there was much interest in analog computing machines. These do not handle numbers in the way that desk calculating machines did then and digital computers do today. In analog machines, numbers were represented by the magnitudes of physical quantities. A slide rule is a good example of an analog computing device. Numbers on a slide rule are represented by distances measured along a scale.

Now, in late 1937, the University of Cambridge took a decision which was very advanced for its period. The university decided to establish a computer laboratory. It was, in fact, called a mathematical laboratory, and it wasn't until around 1970 that we changed the name to computer laboratory—but it was the same laboratory. The university had the intention of equipping this new laboratory with a wide range of desk calculating machines and the latest analog machines, especially a differential analyzer, a machine for solving ordinary differential equations. The results would be presented in the form of a graph, not a table of numbers.

The differential analyzer had been invented some years earlier by Vannevar Bush at the Massachusetts Institute of Technology, M.I.T. We did, in fact, have at that time a model differential analyzer at Cambridge. This was an interesting machine. It was built largely of Meccano parts. Meccano was a constructional toy that was a great favorite with British boys of the period. This model machine was capable of a limited accuracy only and the university proceeded to order a full-scale one to be built according to Bush's drawings, which he very generously made available. This big machine was to be built by an engineering firm. Well, the machine had not been delivered by September 1939 when the beginning of the Second World War brought the university's plans to a temporary halt.

I myself went off immediately on war service, and did not return until September 1945. And then I was appointed head of the computer laboratory. The differential analyzer had been installed, but people's thoughts were already turning to digital machines.

In 1945, there were already one or two digital computers based on mechanical or electromechanical principles in operation, but there was only one electronic machine, one large-scale electronic machine, and that was a machine known as the ENIAC, which had been completed at the School of Electrical Engineering in the University of Pennsylvania in 1945. The program was set up by means of a formidable array of plugs and sockets and switches. The ENIAC was conceived by two very remarkable engineers: Dr. Presper Eckart and Dr. John Mauchly, who were also responsible for its design and construction.

The ENIAC was a straightforward implementation in electronic terms of the sort of machine that a mechanical engineer might have built. It was enormous in size. It contained over 18, 000 vacuum tubes and, as I've said, the program was set up on it by this array of plugs, sockets, and mechanical switches. When the ENIAC was in the final stages of construction and its designers, Eckart and Mauchly, had time to think, they realized that by taking a new and entirely radical approached to the problem, a much smaller and yet more powerful machine could be built.

They had already progressed a good way in this thinking when they were joined by John von Neumann, who was a prominent mathematician of the period. Von Neumann joined as a consultant to the ENIAC project. Now, it's very unfortunate that the results of this highly fruitful collaboration between von Neumann and the engineers was first made known in a report to which von Neumann's name only was attached. In consequence, computers based on the new principles, the principles set out in this report, are often described as von Neumann computers, which does a great injustice to Eckart and Mauchly, who had started the work and contributed some of the fundamental ideas. I myself preferred to use the term “stored-program computer” to describe this new kind of computer. But if names must be attached, I prefer to say the Eckart-von Neumann computer, although I am conscious that still does an injustice to Mauchly.

I first saw a copy of von Neumann's report in May 1945 and it was at once clear to me that this was the way that computer development would go. These were the principles that would determine the structure of future computers. Quite unexpectedly, shortly afterwards I received a telegram from the School in Philadelphia inviting me to attend a course of lectures in which Eckart and Mauchly were to be the principal instructors. At this time, Atlantic shipping was still subject to wartime government regulation and—one didn't go by air then, one still went by sea—I had some difficulty in securing a passage, with the result that the course was more than half over when I arrived. Nevertheless, I learnt, partly from the lectures, partly in conversation with Eckart and Mauchly, all, I repeat all, that was then to be known about the design of stored-program computers. I've always regarded the holding of the School course as an outstanding example of generosity in the sharing of technical information. I have acknowledged it on other occasions. It's a pleasure to me to be able to repeat the acknowledgement in the presence of Dr. Inamori, having received an award from the foundation that bears his name and embodies his ideals.

I stayed on for a few weeks in Philadelphia after the course, visiting people to whom I had received introductions, and began to turn over in my mind the design of a machine that we might possibly construct in Cambridge. And we did, in fact, construct such a machine, and it became known as the EDSAC. Now, the principles were clear, but no stored-program computer had yet been built. To build one was a challenge in electronic engineering and it was a challenge I was well placed to take up. Cambridge was beginning to return to normal after six years of war. There was a great urge there, as elsewhere, to establish peacetime values and to initiate research projects. Moreover, as head of the computer laboratory, I had access to the necessary funds and facilities. And, in addition, my wartime work on radar had added significantly to my experience of electronic engineering. Now, in spite of the great scale of the undertaking, it did not appear that there would be any great difficulty in designing vacuum-tube circuits for performing arithmetic operations—addition, subtraction, multiplication—or in designing circuits for controlling the execution of the program. And I'm sure that other electronic engineers setting out to design one of the early machines felt the same. Broadly that was true, but, even there, there was something to learn. In many electronic applications, it does not matter if the circuits function incorrectly for a brief space of time. There may be momentary flash on a screen. Probably it passes unnoticed by the viewer. It doesn't matter. But in a computer, the computer would have made a mistake and that mistake would, in all probability, have been fatal. And so to put it in technical terms, we had to learn how to design circuits that would handle transience correctly. Once this problem had been identified and correctly stated, there was no real difficulty in solving it, and I'm sure that later designers of electronic computers had no trouble at all.

So much then for the computing circuits. The digital memory was an entirely different matter, because no digital memory capable of holding hundreds or thousands of numbers had ever existed in the world up to that time. And, of course, the whole project depended on such a memory being available. Now, Eckart, right at the beginning, before von Neumann had joined the project, had made a suggestion for a form of memory that might be constructed. This used an ultrasonic delay unit, and I will explain in a moment how it worked. It was on the basis of this “invention,” if you like, that he and Mauchly proceeded to elaborate their ideas. And von Neumann's report, which he drafted on behalf of the group, was, in essence, based on the same idea.

The delay unit consists of a tube full of mercury about one and a half meters long. At each end there are quartz crystals. Pulses, representing numbers, are fed into the first quartz crystal, which converts them into pulses of very high frequency sound: ultrasonic pulses. And these pulses travel slowly, with the speed of sound, down the tube, taking a millisecond to arrive at the other end. They there impinge on the second crystal, reconvert into electronic pulses, and are amplified and reshaped―they'd become much distorted in their journey. Most importantly they are synchronized. We have a supply of pulses, called clock pulses, coming from a continuously running pulse generator. These clock pulses are fed in; they serve to synchronize the pulses that have emerged from the amplifier and those pulses are put back into the beginning of the delay line. So you see these pulses are trapped. They circulate repeatedly through the delay line and are, in that way, stored. About 500 pulses in the EDSAC design could be stored in a single tank.

Now, this form of memory depended on the known principles of classical physics and there was no reason to doubt its practicability. This was not true of other proposals for a digital memory that had been canvassed at the School and elsewhere. They would all have required research before one could be sure that they were viable.

Now, I was aware that some experimental work had been done during the war years with mercury delay units for purposes connected with radar—not a digital memory—in particular, concerned with the cancellation of permanent echoes. I had no contact with that work myself, but I knew it had been done. And I was extremely fortunate at this point to meet Thomas Gold, who was until recently a professor at Cornell, who had worked at the British Admiralty establishment where this work on mercury delay lines was being done. And Gold was able to give me some essential design information that saved me many weeks of experimental work.

I figured out that we would need 32 of these delay tubes—they were called tanks, for some reason we called them tanks—that I would need 32 tanks and I decided to build them in batteries of 16. You can see on the top there are six tubes; underneath there are two other layers and they are connected to electronic equipment. There were very stringent requirements on the mechanical construction. The crystals in the various tanks had to be the same distance apart, to within a small tolerance. That is, the length of each tank had to be the same to within a small tolerance. And, moreover, the crystals had to be accurately parallel to one another. Now, I had full confidence in my calculations—that this would be alright—but nevertheless, I confess that when the first battery was constructed and made to work, I experienced a feeling of relief.

The construction of the EDSAC proceeded from that point with great speed. Early parts were constructed while later parts were still being designed. And much of the load of making the machine work fell on the shoulders of William Renwick, who had joined the project at an early stage as principal engineer. Quite suddenly on the 6th of May, 1949, the EDSAC did its first correct calculation. The EDSAC is very much smaller than the ENIAC—it only has a little over 3,000 vacuum tubes, instead of six times that many—and yet nevertheless it was a much more powerful machine. The power arising from the fact that it had been designed on the new logical principles.

At this point, the machine was working and the second stage of the project could begin. That is to say, the development of programming methods and their application in as many scientific fields as possible. David Wheeler, a research student who joined before the machine was working, contributed greatly to the development of programming methods and his work has been widely recognized. Towards the end of 1950, we were in a position to write a comprehensive report on the methods that we had developed. This report was published in the following year, 1951, in the United States by Addison Wesley under the joint authorship of myself, Wheeler, and Gill, another research student who had likewise made major contributions. This book was entitled The Preparation of Programs for an Automatic Digital Computer with Special Reference to the EDSAC and the Use of a Library of Subroutines. It was, it fact, the very first book on computer programming that was ever published. A copy of this book found its way to Japan, and Dr. Goto and Dr. Wada have told me that it influenced their early thinking.

Now that the machine was working, I naturally began to think about what improvements one could make to the method of design of a computer. I was very much concerned with maintenance, because the vacuum-tube circuits were not very reliable—nothing like the reliability of modern semiconductor equipment. They required a great deal of maintenance. It was clear to me that those sections of the machine, for example, the memory of the EDSAC, that were regular in construction—that is to say, they were built of replicated, identical units—that they were much easier not only to design, but also to maintain, than those parts of the machine which did not have that regular structure. The arithmetic unit of the EDSAC, where the computing was done, was not at all regular, but one could see that there were other ways of designing an arithmetic unit that would be regular. But the control presented a problem. I thought about this for a long time and finally I proposed a system that I called microprogramming. I regarded the control circuits of a computer as a kind of little computer in miniature, a little computer inside the big one, and this little computer would have all the flexibility of a programmed computer.

Our second machine, EDSAC 2, was also a vacuum-tube machine and it was based on the use of this principle of microprogramming, and was working in 1958. There is no doubt that it was EDSAC 2 that established the viability of microprogramming, although vacuum-tube technology did not lend itself at all well to the design of a microprogramming control unit. This situation changed completely with the coming of transistors, and, as a result, it was not until transistors came into general use that microprogramming itself become widely adopted.

I've not time to describe in detail the many other things on which I worked with my colleagues and students in subsequent years. They included further developments in programming methodology, early developments in time sharing, use of cash memories, machine-independent compiling, things like that.

I might mention one development in the years immediately before I left the laboratory—I left in 1980 and went to work in the United States. It was in wide-band, local-area digital communication. It became clear to me, and indeed to others at that period, that progress in the communications associated with the computer depended on breaking away from the traditional techniques that were used in telecommunications and using techniques based on those we had developed for the computers themselves. We built a digital ring, which was a great success. The development of that ring coincided in time with the development of the Ethernet, whose designers had similarly appreciated the need to go to computer techniques and they produced a very elegant system, called the Ethernet, which was quite different from the Cambridge ring but its role and purpose were same.

I've been asked to say in conclusion something by way of advice to young people who hope to achieve distinction in technology or some other field. But, in fact, apart from general advice like work hard and don't get distracted, and so on, there is really very little that I can say, for this reason: that the world changes to a very significant extent during a person's lifetime. I remember my father used to say in his later years how much he wished that the world had remained as it was until he retired. And many other people have felt the same.

It is, of course, true that as one gets older, one is conscious of having learnt many things and gained much experience. But I think it would be useless for me to pass this experience on to you—and I'm now speaking to the younger people present—because the advice that I would be giving you would be advice that applied to the world that I have known, not to the world you are living in. And so I would advise you, in all seriousness, to listen carefully to what older people have to say, but not to allow your natural respect for them to influence you unduly if it is clear that their remarks are not really relevant to your own situation.

We are, of course, all limited by the opportunities that present themselves to us. Dr. Amos Joel, who received a Kyoto Prize in 1989, said in his commemorative lecture that as a young boy he developed a strong desire to work in telephone switching. By very good fortune, his father knew a neighbor who was an executive at the Bell

Telephone Laboratories and this friend was able to give the young Amos Joel the opportunity that he required. And similarly, I was fortunate to become head of the computer laboratory when digital computers were new. If Dr. Joel and I had not had the good fortune we did have, it is possible that we would have received Kyoto Prizes in other fields, but one can, by no means, be sure. A person's career depends, to a major extent, on the opportunities that present themselves to him. Thank you.