An interview with

James H. Wilkinson

Conducted by J. C. Nash
on
13 July, 1984
University of Waterloo, Waterloo, Ontario, Canada

with permission of J. C. Nash

ABSTRACT

James H. Wilkinson recalls his early interest and involvement in scientific computing. He recounts his pre- electronic computer calculations at the Cambridge Maths Laboratory and at the National Physical Laboratory in Teddington on applications such as supersonic flow, thermodynamics and ballistics and the importance of getting mathematicians to take numerical analysis seriously. He discusses his postwar work with Alan Turing on the design and programming of the Pilot-ACE and how it informed his work in round off error analysis and solving eigenvalue problems.

This interview took place on July 13, 1984 at the Ninth Householder Gatlinburg Conference held at the University of Waterloo, Waterloo, Ontario, Canada. The interviewer is J.C. Nash.

NASH: How did you get involved in scientific computing?

WILKINSON: I always liked arithmetic. I was not one of those mathematicians who dislike numbers. I really enjoyed computing with pencil and paper. I well remember learning of Horner's method at school and immediately using it to find the real root of a sextic [a polynomial in x^6] to several decimals without a computer of any kind. Yes, it is fair to say that I had a natural liking for arithmetic.

Then I took an Honours Maths course at Cambridge and in three brief years that enthusiasm was almost completely extinguished.

NASH: With Hardy? [G.H. Hardy, a pure mathematician famous for abhorring applications]

WILKINSON: Yes, I knew him quite well, oddly enough, mainly via cricket. I met him at a cricket match and we got acquainted. He wasn't, of course, a great personal friend of mine; remember, I was a young lad and he was an internationally famous mathematician. If at the end of those three years somebody had forecast that I was going to spend the greater part of my life in the computing business, I would have told him he was crazy.

Then in 1939 the war came. I had either to go into a government department or into the infantry and that was a fairly easy choice to make. The government department happened to be in Cambridge, so I had the good fortune to stay in Cambridge for another four years. I worked in what was known then as the Maths Laboratory and our assignment was to solve day-to-day problems arising in connection with the war effort. So I worked on a variety of rather shaggy and untidily formulated problems. Naturally we would first try to solve them by analytical techniques and when that failed (as it usually did), we would try to solve them by numerical methods on desk computers. Gradually the tide turned and my interest in computing revived.

NASH: 'Desk computers' meaning hand cranked...

WILKINSON: Hand cranked computers. We did have some electrical machines: Fridens and Marchants, but there weren't many of them and I was not fortunate enough to get one. However, we soon got very skilled at using the hand computers; I particularly liked the Brunsviga and became so familiar with it that I doubt whether I could have been faster on a Friden. Most of my war-time computing was done on a Brunsviga, though later at the NPL [Ed. U.K. National Physical Laboratory, Teddington] I used a Facit hand

computer, a beautifully designed machine. As I said, I worked on a variety of problems. I worked on supersonic flow, the thermodynamics of explosives, internal and external ballistics and fragmentation.

NASH: What does 'internal ballistics' mean?

WILKINSON: Internal ballistics is the study of the reactions which go on inside a gun; it includes the study of the detonation wave,the motion of the shot and the recoil of the gun. Fragmentation is concerned with the break-up of an exploding shell and is mainly statistical. But gradually the main mathematical problem of interest to me became the solution of partial differential equations associated with supersonic flow, and I solved one or two such problems with curved shock waves on a desk computer, a very laborious occupation.

NASH: With finite-difference methods?

WILKINSON: Yes, and by the method of characteristics which I found very interesting. There was no systematic study of partial differential equations in my Cambridge course. We solved PDE's by separation of variables and even this was done in an ad hoc way. And so here for the first time there was something directly associated with the work which really captured my imagination. You see I had intended to pursue classical analysis and probably would have become a functional analyst in due course, but here was something quite different which I really liked. I read Courant and Hilbert volume 2 in German (with the aid of a dictionary) and was fascinated by it.

The war service lasted for about six years. I was very fortunate in having gifted colleagues. From the rather small department at Cambridge, four of us went on to become Fellows of the Royal Society: Hill, Moran, Sneddon and I. Most of the rest became chief superintendents or returned to the academic world and became professors. So you see, compared with the average government department, the staff was exceptionally highly qualilfied.

NASH: Were you aware at that time of the work of Hartree?

WILKINSON: Oh yes.

NASH: He was doing a lot of calculations.

WILKINSON: We had a differential analyser [DA] at Cambridge which was the next on the production line after the one that was delivered to Manchester (Presumably you know about the Manchester DA?)

NASH: I don't know that one.

WILKINSON: Hartree used the Manchester DA. Ours looked identical to it but since it was in fact the next one on the production line it was a rather nicer model.

NASH: Did that come out of the work of Vannevar Bush?

WILKINSON: Yes, it was Bush's design. Because it was there at Cambridge some people worked on analog side, but I was always on the digital side and the DA and analog computing in general never appealed to me.

NASH: So you moved on from this Cambridge maths group.

WILKINSON: Yes, the Cambridge Maths Lab. was rather a small outstation in the Ministry of Supply. After four years at Cambridge, they decided to assemble a number of such groups together and we were moved down to Fort Halstead. [Ed. near Sevenoaks, Kent]. It really was a fort at the the time of the Napoleonic wars but it was in no sense a fort then. This became the main station for the more theoretical side of research in connection with armaments. In fact, the division in which I worked was called the Division of Theoretical Armaments, if you can imagine such a name. N.F. Mott, who would later win a Nobel prize in physics, was our superintendent. (My wife worked in the same division.) And so I spent the last two years or so at Fort Halstead; it was a rather bigger setting than the Maths Lab. at Cambridge but we were still doing much the same work. However we were in rather closer contact with Service personnel; there were representatives of the Army, Navy, Air Force and problems which trickled from them were often formulated in rather 'popular' terms.

One difficulty was that they did not seem to have much idea of how difficult a problem was. They might ask one to solve a sextic [Ed. the roots of a polynomial of degree 6] one day and be rather suprised that you could solve it because they remembered vaguely that "one couldn't solve polynomials of a degree higher than four". The next day they might ask you to solve four or five simultaneous non-linear, partial differential equations and they would not fully appreciate the disparity between the two tasks. It was a rather chastening experience at times, because failure was much more common than success. Often the problems were just too difficult even by numerical methods.

The nearest equivalent in the USA to Fort Halstead is Aberdeen Proving Grounds. Our reports were all confidential of course so we couldn't publish in standard journals. Reports were automatically forwarded to Aberdeen Proving Grounds. Sometimes we would receive back papers of a theoretical nature and sometimes related experiments had been done at Aberdeen Proving Grounds. In fact I have a rather nice story in that connection with a partial differential equation I solved which was for the detonation of a stick of RDX/TNT in a steel casing. I solved this by the method of characteristics and it has a curved shock wave. My report was duly sent off to Aberdeen Proving Grounds where they experimented with a shell of the dimensions I had treated. In particular they photographed the detonating shell and produced a curve which fitted beautifully on my theoretical curve. However, what they did not know was that I had used the wrong adiabatic in my computation. For a detonating explosive this is a pretty complicated affair because of the very high pressure in the gases behind the detonation wave. I had been provided with the wrong adiabatic.

NASH: This is the boundary value function?

WILKINSON: Yes, that's right. In the meantime I had repeated the calculation with the correct adiabatic and the new theoretical curve did not agree nearly so well. So much for photography! The camera can lie.

NASH: Did you ever figure out why the fit between the wrong calculation and the photograph was so good?

WILKINSON: No, that was never resolved. The episode came right at the end of my career at Fort Halstead. I'm not certain that the corrected computation was ever sent to Aberdeen Proving Grounds. The war had in fact ended although we couldn't get our release immediately.

As I have said, by that time I had become interested in computing again and my experience is instructive in connection with the problem of getting able mathematicians to take numerical analysis seriously. In spite of a natural liking for computing, which many mathematicians never have anyway, a classical education in mathematics had contrived to kill it in about three years. If the war had not lasted so long, it is unlikely that my interest would have been rekindled. You see I HAD to compute whether I liked it or not. I could scarcely sit sulking for six years waiting for the war to end. The interest had returned slowly and if the war had taken only three years no doubt I would have gone back to Cambridge, continued with my classical education and forgotten all about computing. The attractions of numerical analysis are somewhat elusive and you need some experience to appreciate them. This is a central problem we are up against in trying to get good people to take numerical analysis seriously. Their background doesn't make it the natural thing to do, particularly given all the pressures to which they are subjected. I can't imagine Hardy ever encouraging me to solve any problem, however interesting, by numerical techniques.

NASH: So you moved on from there.

WILKINSON: When the war in Europe was over, I began to think to where I was going to go next. I was offered a fellowship at Cambridge.

NASH: You had just the undergraduate degree at that time?

WILKINSON: I never took a Ph.D. [Ed. Dr. Wilkinson took the Sc.D. later, a doctoral level degree based on published works rather than thesis.] A Ph.D. wasn't terribly important in the U.K. at that time. By the end of six years we had reputations such that there was little to be gained by taking a Ph.D. then.

I thought quite seriously of going back to Cambridge, but it was at this stage that I began to hear rumours about electronic computers and these appeared to me to provide facilities for solving partial differential equations in a reasonable way. There was clearly no future

in solving them on hand computers. Shortly afterwards I discovered that a Mathematics Division was being set up at NPL. I got in touch with E.T. Goodwin who had been a colleague of mine at Cambridge in the Maths Lab. He was one of the first to join this new division. He invited me to have a chat with him at NPL and there I met Turing who I knew already by reputation as something of an eccentric. Turing and I had a long discussion and I was very impressed with him. Presumably he must have been reasonably satisfied with me since he said if I came to NPL he would like me to work with him. I think that this offer and my friendship with Goodwin were the decisive factors. So in May '46, six and a half years after I joined the government service, I moved to NPL (as I thought then, temporarily) instead of going back to the University.

Up to that time Turing had worked alone on the logical design of an electronic computer. When I arrived, he had presented his plans to what you might call a "review committee" at NPL. This consisted of a small group of Fellows from the Royal Society. The committee decided that Turing's ideas were basically sound and they gave him a mandate to go ahead and recruit the appropriate staff.

Up to that time everything associated with the project had been done by Turing himself. He was a man with an original and inventive mind. His design had practically nothing in common with the group of computers which arose out of discussions at the Moore School of Electrical Engineering at the University of Pennsylvania. Mauchly and Ekert had already successfully completed the construction of the first electronic computer, the ENIAC (this was not a stored program computer) and their influence was at its peak. The design of the computer EDSAC built at Cambridge by Maurice Wilkes was based on discussions he had at a symposium held there.

I came to NPL in order to work with Turing but, in spite of the fact that he had no other staff, it was agreed that I should work for him half-time. The other half of the time I was to work in what we would now call the Numerical Analysis Section. It was then called the Desk Computing Section and it was the most active section in the division. Among its members, I should mention E.T. Goodwin, L. Fox, F.W.J. Olver and C.W. Clenshaw, all of whom had considerable experience in numerical methods.

However, Turing was not the sort of person for whom one could work half-time. The plan was that I was to learn numerical analysis with Goodwin and Fox since I was in no sense then a numerical analyst. During the war I had merely worked in an ad hoc fashion on a series of day-to-day problems. I was to learn the basic tricks of the trade and after mastering each field I was to try to program the relevant algorithms for Turing's current machine. In the light of the experience gained in trying to do this, Turing and I would continually modify the design. This was quite an assignment; to learn the whole of numerical analysis and design a computer, too. A full-time occupation!

NASH: How long did that stage last ?

WILKINSON: That will emerge. But as you will gather, in the event I worked for Turing virtually full-time. But my dual appointment had one real advantage. Life with Turing

was just one continuous debate and I found it very stimulating. However, from time to time he would have a particularly black mood and I would take the opportunity provided by my dual appointment to spend a few days with Goodwin and Fox and learn some more numerical analysis.

But to return to the computer. When I went to NPL in May '46, Turing was working on what he called Version 5 of the computer though I never saw any documents relating to versions 1 to 4. Turing was not a great documenter and no doubt the earlier versions were buried in the rubble on his desk.

Perhaps I should attempt to give some idea of the flavour of Version 5, a typical Turingesque creation.

It was a serial machine using mercury delay lines for storage, with a pulse repetition rate of what I still call a mega-cycle, being rather old-fashioned in such matters.

NASH: You mean?

WILKINSON: The basic pulse frequency was provided by a master clock which had a one mega-cycle pulse rate. It worked in binary, of course. That decision was taken early on and was regarded as
irrevocable. The word length was 32 binary digits, which is rather better than nine decimals.

NASH: They were fixed point ?

WILKINSON: Yes. They were fixed point but one of the earliest things that I did (at Turing's request) was to program a set of subroutines for doing floating-point arithmetic. These were later to become rather important in the history of NPL. Right from the start Turing was impressed with the importance of speed. It is possibly not widely known that at that time most people weren't. For instance, Wilkes at Cambridge (who quite early became one of our principal competitors) took the view then that electronic computers were so fast that it was much more important to get one built than to make special efforts to increase its speed, and his views were generally shared. Turing took the opposite view and most of the special features of his machine were designed to make it as fast as possible. There was merit in both views but it was certainly true that the machines we were designing then were not nearly so fast as they appeared to be. However, Turing's obsession with speed certainly made for a very untidy machine. A great weakness of mercury delay lines is access time. In order to make them reasonably economic it is necessary to store a number of words in each delay line. Clearly, if one stores consecutive instructions in consecutive positions in a delay line, one could perform only one instruction per major cycle and indeed the early machines (other than ACE) that were based on mercury delays suffered from this weakness. [Ed. ACE - for Automatic Computing Engine - was the name given to Turing's machine.]

NASH: 'Major cycle' meaning?

WILKINSON: Major cycle meaning the time of circulation of the main storage unit(s) each of which held 32 words of 32 binary digits and hence had a circulation time of 1024 microseconds i.e. approximately a milli-second. A conventional design would have meant that the maximum speed of operation was one instruction per milli-second.

NASH: Because unlike a dynamic RAM where you can get at any cell with one or two clock cycles, this had to use a thousand clock cycles.

WILKINSON: The other two early machines to work -- EDSAC at Cambridge (which Wilkes built) and SEAC at the National Bureau of Standards (which Alexander built) -- did, in fact, store consecutive instructions in consecutive positions so that by the time one instruction had been executed the next one had been "missed" and one had to wait a full cycle for it to emerge. To avoid this Turing stored consecutive instructions in such relative positions that the next instruction emerged just when the previous one was completed. Since different instructions took different times for their execution, consecutive instructions were irregularly spaced in the store. As you can well imagine, this made for what one would call "difficult" coding. I'm not sure that "difficult" is the right word. I would say such coding was tiresome or tedious. Also it made the design of automatic programming languages more laborious while at the same time it made them more desirable. However, this feature of the machine turned out to be rather important; it meant we could do up to sixteen instructions per major cycle, i.e. about sixty-four micro-seconds per instruction.

This practice later became known as "optimum coding" or "latency coding", but Turing never used that term. It was characteristic of him to see his machine as the basic one, all the others being out of step.

NASH: The total memory was?

WILKINSON: Well, Turing envisioned a memory of 200 long delay lines, which would have given 6400 words.

NASH: 24 K bytes?

WILKINSON: Yes, and although that may sound rather small now, it was really very ambitious for that time. I am sure Turing would never have contemplated or supported the building of a smaller
machine.

NASH: What actually got built?

WILKINSON: We'll move on to that. Shortly after I joined NPL, Turing moved on to version six and then rapidly to seven and eight. Those were four-address code machines.

[Ed. A machine which had up to 4 address operands after an instruction, one of which would be to give the memory location of the next instruction.]

The earlier machine, version 5, is hard to describe in these terms. But its successors performed instructions of the type A+B to C and selected the position D of the next instruction, which
was necessary because they were not in consecutive positions.

NASH: A complete instruction would occupy one word?

WILKINSON: Yes, but it was a more powerful instruction than that on a conventional one-address code machine. Another striking difference in Turing's design was that he had a number of one-word delay lines and the arithmetic and logical operations were distributed among them. On a conventional one-address code machine the use of one accumulator leads to a tremendous bottleneck. One is always taking one number out of the accumulator to put in another. By having a number of one-word stores (delay lines) this is avoided. You see, this was all related to Turing's objective of making his computer faster.

NASH: And when a word had an instruction, it would also have an address in it?

WILKINSON: Oh, yes.

NASH: So it wasn't like a modern microcomputer with instruction, operand, operand.

WILKINSON: No.

If Turing had stayed at NPL he would have gone for the full scale computer with two hundred delay lines and quite frankly we had neither the facilities nor the experience to embark on such an ambitious project. It should be appreciated that the full-sized computer was far larger than Wilkes was planning, and eventually built as EDSAC.

Although I have used the term 'optimum coding', most programs fell a good deal short of the optimum speed attainable. To achieve this would have been far too tedious. However, when it came to very important subroutines such as floating-point arithmetic, optimum speed was almost achieved. As I mentioned before, I produced the first set of floating-point routines, but when in 1947 Donald Davies, Mike Woodger, Gerald Alway, Billy Curtis, and John Norton joined the team, they all played a part in polishing them up.

NASH: This was all on paper?

WILKINSON: Naturally, we had no working computer. Because of the optimum coding, floating-point arithmetic (and other important routines such as double-length arithmetic) was much faster on Turing's machines than it was on its competitors. The speed of floating-point arithmetic turned out to be very important for me. When we finally built our computer we dusted down our early routines and polished them up further. By the

standards of the time they were very fast indeed and this enabled me to get really extensive working experience with floating-point computation before it was practical elsewhere. I am sure this is why floating-point error analysis first made headway at NPL.

Turing continued with the logical design of machines but after a while he began to get very dissatisfied. The policy had been adopted that the actual construction of the computer should be undertaken by some other government department such as the Ministry of Supply, where personnel experienced in pulse techniques as a result of working on radar were available.

I never liked that decision, but the Director of NPL, Sir Charles Darwin (great-grandson of the great Charles Darwin) was not a very easy man to argue with. Remember, I was quite a junior member of the NPL staff at that time. But as I saw it, there were only two possibilities. Either the external group would be successful, in which case if they had any imagination at all they would take control of the computer themselves. Alternatively they might fail. It seemed to me that we were in a no-win situation and I couldn't understand why Turing accepted the proposal. This attempt to get the machine built outside continued very unsuccessfully and Turing got more and more morose about it.

NASH: This would be about what time, '47?

WILKINSON: Yes, until finally, very belatedly in 1947, Darwin agreed to set up a very small electronics group (not a division) at NPL. It was recruited mainly from people from other divisions of NPL and inevitably most of the recruits were far from being experts in electronics, so they were going to have to learn on the job. A disaster struck almost immediately. The person who was put in charge of the team -- a Dr. Thomas -- is often critized but in my view rather unjustifiably. Thomas was much more interested in industrial electronics than in building a computer. I do not feel that this was unreasonable; it was not easy to have the imagination to foresee that computers were to become one of the most important developments of the century. However, from our point of view Thomas' preferences were unfortunate. But worse was to come. Thomas and Turing had absolutely nothing in common and were scarcely capable of being civil to each other. So there we have the situation where the leaders of the two groups were completely incompatible.

This, naturally, made Turing even more unhappy, and he began to talk seriously of leaving. Finally, he left in 1948 and joined the group led by Freddy Williams and Tom Kilburn at Manchester. They were making rapid strides in the construction of a computer based on what became known as the "Williams-Kilburn store". Turing's decision was, in my opinion, an unfortunate one. He should have returned to Cambridge where he still held a Fellowship at Kings.

I was left in charge of a team which consisted of six people, including myself. We had virtually no contact with the electronics group, and at that stage Goodwin, who was in charge of the Desk Computing Section, had a long discussion with me. He said, "You know this enterprise looks now as though it's going to founder. Before you can be held

responsible for its failure, would you not prefer to become a member of the Desk Computing Section?"

Well, I just couldn't accept that. By this time I was hooked on computers, so I said I would sweat it out and see what could be done.

Then a miracle occurred. Thomas left and went into industry with ICI, where he had always belonged. The person who succeeded him, F.M.Colebrook, was an old radio engineer with very little knowledge of pulse techniques but a great fund of common sense. When he'd been in the post about two weeks, he came over to see me and he said, "You and I appear to be holding a very unhealthy baby." He went on to invite the four senior members of our group to join him in the Electronics Section on a semi-permanent basis and attempt to achieve something together.

NASH: The four of you being?

WILKINSON: Alway, Davies, Woodger, and myself. Woodger later made his reputation in programming languages in connection with Algol. Davies later became the Superintendent of the Electronics Division (as the small electronics section eventually became) and never really assumed wholehearted membership of the Mathematics Division. Alway came back with me to Mathematics.

NASH: So you started off then learning how to deal with the people.

WILKINSON: Yes, learning how to work together and learning electronics at the same time; but remember that most of the electronics group were having to learn about pulse techniques so we were all in the same boat.

NASH: This is when? The beginning of '48?

WILKINSON: This would be about May or June of 1948. Colebrook was a remarkable tactician and soon we were all working rather well together. There were one or two uneasy weeks but soon the animosity died down. E. Newman was in technical charge of the electronics group; he had worked at EMI during the war and already knew quite a lot about pulse techniques. He and I got on remarkably well and that was a great help. In those days supplies were a problem, but fortunately one member of the electronics group, W. Wilson, a giant of a man, knew everybody in the supply world and was able to solve this problem satisfactorily. After we had spent a month or two building bits and pieces and generally finding our feet, Colebrook said "Why don't we get together now and try to build a pilot machine, the success of which will demonstrate to the authorities that we are competent and therefore ensure the continuation of the enterprise". Then in the light of success -- we didn't hint at failure -- we would then go on and build the full scale ACE.

Now it so happened that we had done a little experimental work in 1947 in Mathematics Division when Harry Huskey had spent a sabbatical year with us. At that time we had

designed just such a miniature machine based on Turing's Version 5. This enterprise had been stopped by Darwin when the Electronics section was formed.

To a large extent, we resurrected this machine, incorporating of course a substantial number of improvements. It was to be called the Pilot-ACE and effectively it would be the smallest machine based on the logic of Version 5 which would demonstrate the practicality of it.

NASH: And that would be how large a machine?

WILKINSON: I suppose I was largely responsible for deciding on the size and scope of the machine but any of the other three could by that time equally well have done so. In order to have some specific objective, I decided that it should be capable of solving fully automatically a set of eight to ten linear equations by Gauss elimination. This it would do in a matter of a second or two, very impressive for that time.

NASH: So you needed to store at least a hundred fifty numbers and the word width was thirty-two bits?

WILKINSON: In fact we decided to have ten long delay lines, that is 320 words. We started to design the chassis in late '48, some chassis being designed by the "mathematicians" and some by the "engineers". In the event, the mathematicians probably designed slightly more than half the chassis. I must emphasize that I am now talking about the detailed electronic design, not just the logical design. We put our newly won knowledge of electronics to immediate use.

NASH: How long did the construction phase take?

WILKINSON: We started to send our blueprints to the NPL workshop towards the end of that year. As each chassis arrived from the workshop, we put it into the main frame.

NASH: Literally a main frame.

WILKINSON: Yes, there really was a frame. We decided to use a plug-in assembly and planned to have spares of key chassis.

By the standards of the time it was an incredibly small machine physically, and yet it was in many regards more powerful than either EDSAC or SEAC. Direct comparisons are not really possible but Pilot-ACE was substantially faster on most problems and it could solve some problems the other two couldn't.

NASH: And the clock cycle was still one mega-cycle?

WILKINSON: Yes, still one mega-cycle, a slightly tough decision.

Wilkes had decided on 500 Kc. Certainly some of the problems we had would have been a lot easier at 500 Kc's. [Ed. Kc = kilocycle = KHz].

NASH: It is interesting that the Apple II is a one mega-cycle or one mega hertz machine, by comparison. [Ed. I was thinking here of the instruction rather than clock rate.]

WILKINSON: Yes, that's right.

The completed chassis would have started to arrive, I imagine, well through '49;,I'm afraid progress was not documented. It so happened that the first chassis to arrive had been designed by Alway and myself, two of the mathematicians of the team, and naturally we put them into the main frame and got them working.

NASH: Did they work more or less the first time?

WILKINSON: Oh no, it took a little while but there were no major design faults.

NASH: You didn't have a lot of patches to do.

WILKINSON: No, on the whole it went pretty well.

Then when the next chassis arrived -- which Alway and I had not designed -- we assisted in its installation because we already knew about the earlier chassis. Thus, without any conscious decision being taken, Alway and I became the debuggers.

NASH: Weren't the chassis somewhat different from each group? Or were these different components?

WILKINSON: Of course, the various chassis had entirely different functions. Thus several were associated with the line counter, several with the logical control and then there was one chassis for each delay line. (The latter were, of course, all identical.)

NASH: The line-counter is?

WILKINSON: This was the section which counted the basic 32 pulses in a word time.

NASH: All this is now on one chip?

WILKINSON: Yes, of course, and much more. Our units were vast by today's standards in spite of being small by the standards of the contemporary design. Pilot-ACE was also unique among the early computers in being extremely mobile. The main frame was on wheels and when the computer was finished, we wheeled it back to Maths Division without affecting its performance.

NASH: Was it power hungry?

WILKINSON: It consumed somewhat less than ten kilowatts, which was quite low.

NASH: Yes, that's quite low.

WILKINSON: But we didn't have any forced cooling and perhaps the construction was a little too compact for that. When we were assembling it, we were, of course, standing in front of it all day. It was like working in front of a 10 kilowatt fire, a rather trying experience.

NASH: It was generating about the same energy as ten electric kettles.

WILKINSON: Ten kettles spread over an area about three feet by fifteen feet.

NASH: Did you have much component trouble?

WILKINSON: Not really. Our main problem was with germanium diodes, which fortunately we didn't use on the scale as SEAC. We used them for some gating requirements. Because our machine was so compact and didn't have forced cooling the diodes were working at a temperature which was much higher than specified by the manufacturer.

NASH: So they would fail?

WILKINSON: Yes, there were some good diodes and some bad diodes. The bad diodes would fail after perhaps a week in the computer.

NASH: Exactly as the experience of today.

WILKINSON: A good diode on the other hand would go on almost indefinitely, so bad diodes were eventually weeded out.

NASH: So you had a literal burn in with a hot chassis.

WILKINSON: Then came a key stage in the assembly of the computer. This was the day the first delay line chassis was integrated. This was designed by Newman and as usual he joined Alway and myself while it was installed but from then on he stayed with us. The three of us worked well together, and debugged the whole of the rest of the machine.

NASH: When did the first program run?

WILKINSON: On May the tenth, 1950. It is interesting that unlike Wilkes, who had built everything he intended to have and then made it work, we added chassis by chassis as they were completed and as soon as it was possible to do something (which was as soon as we had the control unit working, the adder and the subtracter, the logical operations and one long delay line), we tried it.

NASH: How would you feed the data in?

WILKINSON: Oh, at that point we fed the instructions in (in binary) from a set of 32 keys. When it worked on May the tenth, it would could perform only the simplest of programs. In fact our first program achieved the following: it took the binary number set up on the 32 keys and every major cycle it added that number into the accumulator until it overflowed. Now, in addition to the 32 input keys we had a set of 32 output lights. When an overflow took place, the program put on the next light. So successive lights would come on at a speed which was directly related to the size of the number on the switches. Now this program, admittedly rather small, had to be fed in one instruction at a time, in binary from the 32 keys. At the time the design of the delay lines needed improving; the amplifiers were somewhat unstable. So we kept feeding in the program and it kept being forgotten before we could complete the input. So I said to Alway:"Let's try it four or five more times and if it doesn't work, we'll call it a day and go home."

Well, we put it in about four times and suddenly all the lights came on. Well, this could have happened in any case, and it didn't guarantee the program was working. However, we made the input number smaller and the lights came on more slowly.

NASH: So the amplifiers had settled down?

WILKINSON: Yes. Then we doubled the number, and the lights came up twice as fast. We made the number three times as large and they came up three times as fast. On a binary machine that was quite convincing, so we said:"It must be working" and went home rejoicing. That program later became rather famous on the machine. It was known affectionately as "Successive Digits" or SUCK DIGS.

Sometime before this , Teddy Bullard (later Sir Edward) had succeeded Darwin, and when he visited the Electronics Section (in late April, 1950) he asked me how it was going. I replied that we should have something going in a week or two. Bullard was a very forthright chap, and he said with some scorn "Come on, you can't pull the wool over my eyes. I've heard it's going very badly." (He had heard this, quite justifiably, via Harry Huskey.) I said "You may well have heard this and indeed it was true, but it's coming along nicely now and in a week or two, I confidently expect it to be working."

NASH: So you were right on target.

WILKINSON: Naturally, when it did work, I tried to get in touch with him as I had promised to do. I tried to 'phone him. He wasn't there. Now the machine wasn't really very good at that stage and might stop working at any time. The director could not be traced and I was pacing up and down, saying:"The bloody director is never here when you want him" when he stepped into the room via the window. His opening words were "Here's the bloody director. I hear it's working."

I showed him this program and he played with it and agreed that it was working. Then he turned to me with a grin and said "It may be working, but the program's somewhat less than epoch-making," with which we had to agree, but it was very heartening for us.

We continued to add the chassis one by one, and by the end of June most of it was assembled. We didn't at that time have a multiplier, nor had we planned to have one, on Pilot-ACE.

NASH: You would use successive addition?

WILKINSON: Yes, it was to be done by a subroutine, optimum coded so that it was not too slow. In fact, the optimum coded version was about as fast as the automatic multiplier on EDSAC. So as soon as it began to do significant things, Bullard began to press us to have an Open Day and to demonstrate it to the world. Well, I was a bit anxious about that because it wasn't really reliable enough. The amplifiers on the delay lines were still inclined to be unstable. However, Bullard was a very impetuous man and he finally landed us with these 'demonstration days'.

NASH: When was that?

WILKINSON: It would have been November of 1950. By that time we could do a variety of significant things, but it was still not a very reliable machine.

NASH: Was a cool day more helpful than a hot day? And a dry day better than a wet one?

WILKINSON: One of the troubles we had at that time was with the power supply. Not our power supplies but that of the Central Electricity Generating Board. For instance, in the evening when everyone arrived home and switched on electric fires, the voltage would drop suddenly, and that gave us problems.

NASH: Historically there was a coal shortage at that time?

WILKINSON: Yes. Such things added to our difficulties. We knew, too, that when SEAC had had its first demonstration - a little before us - it had been a fiasco, even though SEAC had, in general, been working reasonably well. During the whole of the time allotted to the press demonstration, it never once worked. You will find the early years abound with such bad luck stories.

I must confess to having been pessimistic. We decided to have two popular programs for the daily press. For the first, they would give us a 6 figure decimal number and the computer would tell them if it were a prime, and if not, output a factor.

For the second program, they would give us any date from the year zero up to the year 9999 and it would output what day of the week it was. It covered both the Julian and Gregorian calendars and dealt with all leap years. In all, quite an amusing little program.

Mike Woodger produced that program.

NASH: And where did he discover the technique?

WILKINSON: He worked it out for himself. Such programs are good fun of course, but they leave one mercilessly exposed to the vulgar gaze. Someone puts in the current date which is Wednesday,say, and the machine promptly says Thursday! So they're very much more dangerous. If you tell the press it's solving a partial differential equation, you can swear blind it's solving a partial differential equation and they would be hard put to prove it is not. Finally, we were to have one serious program; this traced skew rays through a set of lenses.

Well, we decided on this last program and announced it only to find that we couldn't get the program to work. Two days before the press show it had still never worked, and we didn't know whether the program had a bug, or whether it was due to computer malfunction. Then, just two days before the show, Alway and I accidentally found it was a minor machine fault which was not invoked at all in our other programs.

We got all three programs working then, just in time. The arrangement was that Bullard would entertain the popular press and I would give the demonstrations. The whole thing was to cover three days; one day with the popular press, one with the technical press, and a third day for VIPs including our competitors. Wilkes had his machine running in Cambridge and was justifiably proud of it. Williams and Kilburn from Manchester were also coming.

NASH: They had a machine too, didn't they?

WILKINSON: They had a little hook-up at that time but it could scarcely be called a computer. They hadn't built the Mark I by that time. Their little hook-up was the first anywhere ever to run a stored program. It worked in 1947 and found the highest common factor of two numbers. This was, of course, a great deal smaller even than the Pilot-ACE. However, it WAS an impressive 'first' and I well remember being very heartened when I saw it working.

My point then, is that Open Day was doomed to be a failure. The plan for the first day was that Bullard was to entertain the press upstairs, while downstairs we made sure the computer was working. We were to receive a signal when Bullard was almost through. We did, and immediately the machine stopped working. We found out, almost at once, that it was a chassis associated with one of the delay lines. We plugged in a spare, but unfortunately we knew that the amplifier, as it warmed up, would become unstable; the amplifier would then need to be retuned and in ten minutes all would be fine from then on.

So we were expecting to run into trouble almost as soon as the demonstration started. Well, the press arrived. They threw numbers at us and the computer factorized them like a charm. It was indefatigable!

We moved on to the 'dates' program. It worked as it had never worked before: the day of Trafalgar, Waterloo, King George V's birthday.

NASH: You were very lucky.

WILKINSON: We moved on to the ray tracing. It traced rays like a fiend; nothing could stop it. It continued in this vein from ten till one o'clock. Then the press went away to lunch. We immediately looked at the output from the delay line, that is the shape of the pulse coming out. It was the best output we'd ever seen! Further press representatives came in the afternoon; still a faultless performance.

The next day we had the technical press and it was the same story. Never before had it worked for anything approaching this time period without a fault. The third day the VIPs came. Surely it would let us down now? Not a bit of it. Wilkes was there. I have always found him a very fair man, but naturally he was not prepared to give anything away. He didn't get a chance; it was perfect. It had already been decided that there would be a fourth day when it would be put on show for the staff of NPL. This was a Saturday. The computer had a small fault before our audience arrived but we soon put this right and again it performed flawlessly. The chances of such a performance must have been a million to one against.

On Monday we came in feeling rather jubilant. The computer was down and it took us about a week to get it working again!

NASH: What did it have as peripherals at that time?

WILKINSON: By that time the input and output was by Hollerith punched cards. We never had tapes.

NASH: And would these be hand-fed or machine-fed cards?

WILKINSON: Machine-fed. The reader fed 200 cards per minute.

NASH: Did you have to make that reader?

WILKINSON: We had a Hollerith section in Maths Division, so that there was quite a lot of expertise in that sector. In fact this side of the work was done by Stan Gill who was then at NPL but was later better known for his work at Cambridge with Wilkes and Wheeler.

NASH: Using the IBM-type tabulator?

WILKINSON: It was the card reader which read at 200 cards per minute; the output card punch operated at 100 cards per minute. We never had any direct printout from Pilot-ACE.

NASH: You would run the cards through a listing machine?

WILKINSON: Yes, that's right.

Now the Pilot-ACE had been conceived in an atmosphere of crisis. There was the danger that if we did not make some progress the project might be closed down. I had never really formulated any clear idea what we were going to do next but it had never been my intention that we would use Pilot-ACE for serious computation. It suddenly dawned on me after the fantastic success of the demonstration that we were expected to take the computer back to Mathematics Division and do day-to-day computing on it. You know, I hadn't thought of it in those terms; I had always thought that when it was successful the authorities would say:"Good. It works very nicely. Now go ahead and build the full-scale ACE." But, of course, that was not possible now. So we embarked upon a short program of improvements to the computer, including the addition of an automatic multiplier, and then we took it back to Maths Division. The Pilot-ACE was used up to mid-1955, so for four and a half years. It worked pretty well by the standards of computers at that time; rather better than its competitors.

NASH: Programmed in machine code?

WILKINSON: Programmed always in machine code. We didn't start writing programming languages until later. In 1951 English Electric decided it wanted to get into computers and agreed to start by copying a working computer. So it copied Pilot-ACE with some improvements and called it DEUCE. This was effectively an engineered version of Pilot-ACE; it was marketed and sold over 30 copies which was then considered to be a substantial number.

NASH: Was there a full ACE?

WILKINSON: Yes, we later built a full-scale ACE at NPL, and despite what has been written about Turing's plans, the full ACE realized everything that Turing had in his designs and a good deal more. However, in retrospect, the construction of ACE can be seen to have been a mistake. It was completed in 1958. By that time new techniques had been developed. Alway, Woodger and myself played no part in the construction of ACE (we were back in Maths Division) but I do not imply that this was the reason for its comparative lack of success.

NASH: It was still using mercury delay lines [for memory]?

WILKINSON: It was. This in itself was a bad decision, I think. However, it was an interesting machine and fun to use. Had it been completed in 1956 it might have been a different story.

NASH: So you used Pilot-ACE. Did you use DEUCE?

WILKINSON: Yes. NPL had the first DEUCE to be built. So we used Pilot-ACE up to mid '55. Then the DEUCE arrived.

NASH: The DEUCE was a little bigger?

WILKINSON: Yes it was a little bigger (physically it was a lot bigger and therefore did not get so hot) but the principal change was the addition of a drum drive. After we came back to the Maths Division the Electronics section worked on the development of a drum and this was added to Pilot-ACE after we had been using it for about a year. The first drum stored 8096 words. The transfer of information from the drum to the so-called high-speed store was completely autonomous so that computing could continue while the transfer was taking place. This meant, for instance, that for matrix work, we could work just about as fast as if we had the full number of delay lines [Ed. as in the full ACE design] because the transfers were always done while the previous data were being used.

NASH: And did you ever add direct printing?

WILKINSON: No, never. Neither Pilot-ACE nor DEUCE, or for that matter on ACE. The only output throughout their lives was by punched cards.

NASH: DEUCE lasted until about '62?

WILKINSON: As a matter of fact, the NPL DEUCE continued to be used long after we had the full-scale ACE. We had numerous long-term visitors and many of them had written programs for DEUCE which they had cleaned up rather well. They were reluctant to re-write them; it's an old story.

NASH: And all these machines had the same machine code?

WILKINSON: There were slight changes in the light of experience. We realized there were some things we could do even within the basic limitations of the Pilot-ACE which would make it much more flexible to use and those were incorporated in DEUCE.

NASH: And DEUCE had programming languages?

WILKINSON: We began to develop programming languages on it.

NASH: An assembler?

WILKINSON: Oh yes. And there was a rather nice matrix interpretive scheme, not very sophisticated, but very effective. It enabled you to do the sort of thing that Cleve is demonstrating here but not, of course, in so convenient a way. [Ed. Prof. Cleve Moler of the University of New Mexico, whose MATLAB will hopefully be discussed in a later

issue.] One could build up quite complicated matrix operators from standard "bricks" as we used to call them. It was much nicer than using subroutines.

NASH: Interpreted?

WILKINSON: Yes.

NASH: The ACE and DEUCE work led to two books from you -- Rounding Errors in Algebraic Processes, in 1963, and later the Algebraic Eigenvalue Problem, in 1965.

WILKINSON: Well, that's rather a long story. In spite of popular belief, I was not born with a passionate interest in linear algebra or rounding errors.

NASH: But it's what you're best known for.

WILKINSON: As I've mentioned earlier on, I intended to work on partial differential equations and was itching to get back to them. One of the earliest things I did was to rework the problem on RDX/TNT I had done at Fort Halstead and I was delighted to find that the results (using the right adiabatic!) were correct. But then, to my immense surprise, we were deluged with a succession of problems in the numerical linear algebra. The aircraft companies were interested in building bigger planes and this led to much bigger free vibration problems. They also needed to solve the flutter problem in many more variables than hitherto. All this work gave rise to matrix problems. The director of NPL, Bullard, was interested in the origin of the Earth's magnetic field. He'd reduced this, by the time he came to me, to an algebraic eigenvalue problem. I also did computing for Aldermaston [Ed. UK Atomic Weapons Research Centre], and rather surprisingly the work they asked me to do was also in the linear algebra field. Because of its peculiar design features, Pilot-ACE was exceptionally efficient in that area relative to other computers.

NASH: It had double precision accumulation of inner products?

WILKINSON: Yes. Double precision accumulation was very efficient on all ACE-type machines. Even before we added the drum storage, the way Turing had organized the use of the card reader made it possible for us to do unsymmetric eigenvalue problems up to order 60 in spite of the small memory.

NASH: By feeding information out and back in?

WILKINSON: Not exactly. This was achieved as follows. The matrix was stored in binary with two numbers on each row of a card. To perform an iteration the cards would be read at top speed and all the computation was done between the rows of the cards. There was about 15 milliseconds between the rows of the cards, and to deal with two elements of the matrix took 12 msecs.

NASH: So you were computing as it was reading?

WILKINSON: Yes. The net effect was that we could work almost as fast as we would have done if the data had all been in our mercury delay lines. In fact, if EDSAC had had enough storage to hold the full matrix, it still could not have computed as fast as we were using the cards. This was all part of Turing's design. It was a rather ragged, slightly untidy machine, but a very interesting one, and its design certainly owed very little to others.

NASH: So when did you actually start writing "Rounding Errors..."?

WILKINSON: That is rather a long story. As a result of actually solving a large number of matrix problems, I bacame aware that the algorithms I was using were much more stable than everybody had expected. Naturally I began to wonder why. This started me on rounding error analysis. The first analyses I did were in connection with polynomials. I moved on to analysis in the field of linear algebra around '55-'56, but I didn't start to publish my work until '58-'59 . You see, I hadn't acquired the habit of submitting my work to the standard journals. During the war we weren't allowed to publish and then I had spent several years designing a machine and then building it but I had no intention of making my career in those fields. Then when the computer was built we were so hard-pressed using it that I didn't write things up. So if you look at my publication record you'll find amazingly little published before 1957, and then a sudden outburst. I was really writing up what I had been doing in the last four or five years.

NASH: The Stationery Office produced "Rounding Errors..."

WILKINSON: Yes, that came out in 1963. I sent off the finished manuscript of the Algebraic Eigenvalue Problem by the beginning of 1964. By the time it had gone through the Press, it came out about March of 1965.

NASH: It actually has the Clarendon imprint, so it went through the delegates and was one of Oxford University Press's flagship books.

Did you move on to KDF9 about that time?

WILKINSON: Yes. Well, we'd always retained our connection with English Electric after they built DEUCE and although the KDF9 design had broken away from the Turing concepts, the ability to accumulate inner products in double precision was retained, fortunately for me.

NASH: Rather than going into that machine, I'd like to ask your opinions on the following. Today a lot of people are coming into computing with no background in calculation. Many of the machines they're using don't have the properties that ACE did, with double precision accumulation of inner products. People have very little knowledge of this. How can these ideas be got across to them?

WILKINSON: It's a really difficult question and I wouldn't claim to know the complete answer to it. Our experience with the Pilot-ACE was really rather special. In order to get the most out of a machine with such a small store, user cooperation was essential on a scale which in many ways is not achieved even now. This gave one an intimacy with the machine; we were forced to look at the numbers and thereby achieved a deep understanding of what was going on. One can, of course, do this with modern computers; indeed the potential for doing it is actually greater, but one has to realize what it is one should be doing and why. For iterative methods we used acceleration techniques which were actually under the direct control of the operator. For instance, when we were using the power method for the determination of the dominant eigenvector of a matrix we could follow the progress of the vector on a cathode ray tube screen. We had a cathode ray display which showed the contents of any long delay line. We would look at the screen (which showed 32 components of the current vector) and we would see how fast it was converging. We would put a piece of paper over it and we could say, for example, "It's gaining a binary digit every three iterations so the ratio of the dominant to the subdominant eigenvalue must be about 2 to the power 1/3 ." We could then set up a shift of origin on the input keys that would give much faster convergence. This work was commonly done by assistants who were in no sense qualified mathematicians but they became very expert indeed. It is surprising how well they understood the battery of acceleration techniques available and how efficiently they used them. When later we went over to more automatic techniques they complained we were "taking the guts out of their work". They really loved these early programs.

The familiarity and intimacy gained with the computing process was fully comparable with that which one gets on a hand desk machine where perforce you see every number. But on ACE that familiarity was gained quickly and painlessly. This experience was invaluable. Is there any way you can get it now? Of course there is, but one needs to know what is worth having and to have the incentive to output it.

NASH: I would say my own experience is that we are transferring large machine, "faceless" programs down to the personal computers where in fact one can go back to the ACE/DEUCE ideas.

WILKINSON: Yes. I agree. The potential is there and it's much greater really than it was on ACE/DEUCE. But in my experience, many people who do computing are reluctant to look at numbers. At Stanford the general level of our students has been pretty high but I would say their main weakness is in their inability to look at outputs and extract the meaningful information there is in them. In fact, somewhat to my surprise, they are generally less efficient at this than the assistants I used to have at NPL in the ACE/DEUCE days, in spite of having far superior mathematical qualifications. Most of those assistants had experience with desk computers and had learned to "look at numbers". The Pilot-ACE forced them to continue with this habit.

I certainly do not want to suggest that the way to acquire this habit is to serve an apprenticeship on hand desk computers, but we have yet to learn how to instill the relevant knowledge.