

Tom Kilburn: A tale of five computers

One of the pre-eminent figures in the early history of computer design was Tom Kilburn. Over the course of some thirty years, he made significant contributions to the development of five important computers. Although a natural team leader possessed of a somewhat dominating personality, who inspired in those who worked closely with him great loyalty and affection, Kilburn was, on casual acquaintance, a self-contained man who chose his words with care. As F.C. Williams put it “What you must always remember is that Tom is a Yorkshireman.”¹

Early Days

Tom Kilburn was born on the 11th August 1921, near Dewsbury in West Yorkshire, England. His father, John William Kilburn, was a statistical clerk who rose to become a company secretary.¹ Tom had a somewhat specialized education at Wheelwright Grammar School having been permitted by his headmaster to study almost nothing else from around the age of 14. It was hardly surprising therefore he emerged from school as something of a mathematical specialist. In 1940, Kilburn went up to Sidney Sussex College, Cambridge, with State, Dewsbury Major, and Minor Open Scholarships. Wartime courses at Cambridge were somewhat truncated and in 1942, Kilburn graduated with First Class Honors in Part I of the Mathematical Tripos and in the preliminary examination for Part II.

During the Second World War, many Cambridge mathematics dons were absent from the university serving at Bletchley Park and elsewhere. In spite of this, there remained a lively mathematical community in which Kilburn played a full part. As the Sidney Sussex college representative in the New Pythagoreans (a subgroup of the Cambridge University Mathematical Society), Kilburn almost certainly came into contact with a number of people who later went on to play a part in the development of computing. Geoff Tootill and Gordon Welchman were, like Kilburn, officers of the New Pythagoreans.² Speakers to the student society included future Bletchley Park code breakers M.H.A. Newman,^{3,4} K.J. Le Couteur, and William (Bill) T. Tutte. However, Kilburn would not have been likely to have come into contact with Alan Turing, who spent most of the years 1937-38 at Princeton based in the world-leading research group in mathematical logic headed by Alonzo Church before taking up his position at for Bletchley Park in 1939. It is also unlikely that Kilburn read “On Computable Numbers” as an undergraduate because his mathematical taste was more applied than pure:

[P]ure mathematics seemed extremely abstract. I was the sort of person who was always prepared to accept that two and two are four, whereas I'd spent the first term at Cambridge in one of Newman's lectures proving that this was so. Whilst it was all very interesting—I mean one could appreciate the beauty of it—it left me rather cold. At the end of it, you didn't seem to be much further forward...⁵

War Service

At some point during his final year at Cambridge, Kilburn attended a talk given by C. P. Snow,⁶ who was visiting universities recruiting people for unspecified war work. By his own account, Kilburn had some fairly settled ideas about what he wanted to do for the war effort:

It seems silly but if I could have joined the RAF as a pilot, I would have done that, but I was relegated to navigator or some such, and that was not quite so appealing. It sounds egotistical but I like to lead. I like to be in charge and I didn't fancy the idea of being driven and crashed by some other character. I wanted to do my own driving and crashing. It's on these sorts of whims that life is founded—it's not through any profound thought is it? You take advantage of what's there at the time.⁵

Having ruled out service in the RAF, Kilburn enrolled on a number of short courses in electronics. He then took a six-week City and Guilds course on electricity, magnetism, and electronics. After about a week-long vacation, he was assigned to the Telecommunications Research Establishment (TRE), Malvern, where he joined Group 19, led by F.C. Williams. He was not, however, greeted with unbridled enthusiasm.

Williams had requested an extra person to join his team, and Kilburn was the person they sent. The other members of Williams' group were all around 30 years old with an average of 10 years practical experience in electronics. By contrast, Kilburn was 21 years old and, prior to being drafted, had not the least interest in electronics or electronic equipment of any kind. Williams, whose group was responsible for designing and debugging electronic circuitry and for solving problems encountered by other groups, made no attempt to hide his disappointment at being offered someone so inexperienced. Kilburn later recalled:

[I]n effect he said "Oh God, you don't know anything?" and I said "No." That was the sort of relationship at the start. But of course by the time we left Malvern—that was four years later—the relationship was quite different.⁵

The Manchester Baby

By the end of the war Kilburn was well settled domestically, having married Irene Marsden in 1943, and professionally, he had, despite such an inauspicious beginning, become an important member of Williams' team. Kilburn had achieved the rank of acting scientific officer. In 1946, Williams left TRE to take up the Edward Stocks Massey Chair of Electro-Technics at the Victoria University of Manchester. It was Williams' intention to continue his work on the development of the cathode ray tube (CRT) memory, and he arranged for Kilburn to work with him at Manchester under secondment from TRE. By the end of 1947, Williams and Kilburn had developed a CRT that could store patterns over long periods. But as Kilburn put it, "the only way to test whether the cathode ray tube system would work in a computer was, in fact, to build a computer."⁷

The story of precisely how the Manchester Baby was conceived, funded, and developed as well as the roles played by various actors in the project is somewhat complicated.^{8,9} The dominant historical narrative has come univocally from the engineering tradition and has generally paid little attention to the contribution made by people like M.H.A. Newman and P.M.S. Blackett.¹⁰

What is not in doubt is that the machine itself was the first working example of a digital electronic stored program computer. However, its significance for the historian of computing is not that it was an iconic first but that it provided the foundation that Manchester used to build itself into a leading center for the emerging computer science field.

Mark I and Mercury

Kilburn's original intention had been to return to TRE after the Manchester Baby was completed, but the success of the Baby was such that the Ministry of Supply quickly awarded a contract to Ferranti to design and build a full-scale commercial computer to Williams' specification. An important initial step was to construct a prototype machine,¹¹ the Manchester Mark I, which was to be produced at the university. It was clear to all that Kilburn was vital to the new project, and Williams, partly with the inducement of a lecturing post, was able to persuade him to remain at the university to work on the prototype.

By the autumn of 1949, the Manchester Mark I, now with backup drum store, was complete, and ran continuously thereafter for almost a year. Between 1951 and 1957 around nine of the Ferranti production versions of the Mark I machines were sold.

Over the three years since the Baby was completed, two important shifts of responsibility had taken place. First, Max Newman came to the conclusion having provided the initial impetus and leadership in the development of the Baby, further computer development required engineers rather than mathematicians to be in the driving seat. Newman's withdrawal left Williams in sole charge, but Williams, who never had much interest in computing as such, fairly quickly passed effective control of further developments to Kilburn.

In 1951, once again following a process of incremental development, Kilburn began working toward a Mark II computer that was known as the megacycle machine, or Meg. It replaced the Mark I valve diodes with solid-state versions and offered a tenfold increase in clock rate together with greatly improved reliability and floating-point operation. The serial CRT memory, which was already running at a near optimal rate in the Mark I, threatened to act as a performance bottleneck for the Meg. Kilburn's solution was to design a 10-bit parallel CRT memory.

Meg first operated successfully in the summer of 1954, and Ferranti developed a commercial version of Meg under the name Mercury. Clients included the Meteorological Office, the Norwegian Defense Research Establishment, and Manchester University. In all, 19 Mercury computers were sold, of which 6 were purchased by overseas customers.

Transistor Computers

Kilburn led a Manchester design team consisting of Dai Edwards and Tommy Thomas, who concentrated on the Meg, together with Dick Grimsdale, and Douglas Webb, who were simultaneously working on what was originally a research project looking into developing the smallest possible economic computer.¹² It was soon clear that a great deal of valuable experience could be gained by using transistors to build the machine. Two prototype transistor computers were commissioned, both of which made use of a pseudo two-address instruction format and permitted optimum programming. The 48-bit machine produced in November 1953, which is widely acknowledged to have been the world's first operational transistor computer, had 550 diodes and 92 point-contact transistors, was manufactured by STC. A somewhat enhanced version of the transistor computer was completed in April 1955, this time boasting some 1300 diodes and 200 point-contact transistors.¹³

In 1956, the Metropolitan Vickers Electrical Company, adapted the design of the experimental transistor computer to allow the use of junction transistors, and manufactured six transistor machines, mainly for internal use, under the name Metrovick 950. From the perspective of Kilburn and his team, the most important aspect of the Transistor computer was the early experience it gave them in transistor circuit techniques.¹⁴

Muse and Atlas

Kilburn's plan with the Muse (microsecond) project was to develop a really large fast machine that would make full use of both existing and emerging technology. He succeeded admirably; Muse used "multiprogramming, job scheduling, spooling, interrupts, pipelining, interleaved storage, autonomous transfer units, virtual storage and paging—though none of these techniques had been invented when the project started in 1956."¹³

Muse can be considered comparable in scope and ambition to the IBM Stretch and Univac LARC projects, but Kilburn was under no illusion that the Department of Electrical Engineering had sufficient resources available to complete, without assistance, a project of this scale and complexity. His initial attempts to elicit Ferranti or the government's support for the Muse proposal were unsuccessful, so the decision was taken to proceed instead with a trimmed-down version of the original scheme. However, in January 1959, Ferranti, with £300,000 in backing from the National Research Development Corporation (NRDC), decided to participate in the project, now renamed Atlas.¹⁴

Among the innovations introduced in the Atlas was a scheme that let programmers treat drum stores as if they were core storage. An innovative program called the *supervisor* managed drum transfers. The *one-level-store concept*—the idea of a fast and a slow store appearing as a single fast store—was an important precursor to virtual memory.

Kilburn was not only responsible for the management the project but also played a part in the circuit design, including work on an adder with a fast carry path.¹ In general though, he relied substantially on the experienced teams that he had established over a number of years. Three Atlas systems were eventually built and installed at the Universities of Manchester and London and at the Rutherford Laboratory.

In 1960, Kilburn was appointed a professor of computer engineering in the Department of Electrical Engineering.

Department of Computer Science

In addition to his considerable contribution to early computer design, Kilburn also did much to establish computing as an academic discipline in the UK Higher Education curriculum. Starting in 1963, he spent several years establishing and organizing a new Department of Computer Science, the first of its kind in the

UK. The intention was to provide a natural home for computer research, a sound base for future projects, and undergraduate courses in computer science. Kilburn, now translated into a professor of computer science, was the inaugural head of department, and had 12 academic staff under him. Directly reflecting Kilburn's personal strengths and his professional experience, Manchester placed more emphasis on hardware than many of the other computer science departments that followed it, most of which sprang from a mathematical lineage rather than from engineering. Kilburn went on to serve the dean of the Faculty of Science from 1970 to 1972 and pro vice chancellor between 1976 and 1979.

MU5

In 1966, Kilburn embarked on what was to be his last major computing project: the MU5. The Atlas had been operational for four years, and the MU5's main focus was to provide a computer architecture geared toward the efficient running of programs written in high-level languages. The MU5 was conceived as a range of three machines—a small inexpensive computer, a high-spec scientific computer with 20 times the throughput of the Atlas, and a multiprocessor—but only the second was actually developed.¹⁵ The MU5's original design proposal was set out in 1968 at the Edinburgh International Federation for Information Processing (IFIP) conference in a paper authored jointly by Kilburn, Derrick Morris, Jeff Rohl, and Frank Sumner.¹⁶

An interesting technical aspect of the MU5 was the associative name store in which frequently used scalar variables would automatically reside in a fast cache store. Morris explained, "This was as a result of an analysis of the Atlas software, especially the instruction code. We learnt something about the frequency of use of operands and control structures. The order code accommodated string functions and vector functions."¹⁵

The university secured the cooperation of International Computers and Tabulators (ICT), which made construction facilities available at cost and provided five staff to work on the project. The university's relationship with ICT and the potential it created to benefit the company persuaded the Science Research Council (SRC) to assist the project by awarding the university a £630,000 grant over a five-year period.

It was a fruitful collaboration. However, an initial failure by International Computers Limited (ICL) (which by then had merged with ICT) to acknowledge the extent to which the MU5 had influenced their 2900 series concerned the SRC, outraged Kilburn, and led to a long-running dispute that was not fully settled until after Kilburn's retirement in 1981.

Retirement

In order to spend more time with his wife, Kilburn retired early at age 60. Unfortunately, Irene Kilburn died just two weeks before his planned retirement.¹⁷ After that, he continued to spend one day each month in his old department, but the majority of his time was spent with his son and daughter, gardening, playing the piano, and following the Manchester United Football Club. He died in Manchester on 17 January 2001.

References and notes

1. M. Wilkes and H.J. Kahn, "Tom Kilburn CBE FEng. 11 August 1921 – 17 January 2001," *Biographical Memoirs of the Fellows of the Royal Society*, The Royal Soc. London, 2003, pp. 283–297.
2. Geoff Tootill was Christ's College representative for the New Pythagoreans and president of the Archimedean, and Gordon Welchman was a student at Sydney Sussex College and honorary vice president of the New Pythagoreans.
3. D.P. Anderson, "Max Newman: Forgotten Man of Early British Computing" *Communications of the ACM*, Vol. 56 No. 5, Pages 29-31
4. Kilburn and Tootill were also students in some of Newman's classes.
5. G. Bowker and R. Giordano, "Interview with Tom Kilburn," *IEEE Annals of the History of Computing*, vol. 15, no. 5, 1993, pp. 17–32.

6. Originally trained as a chemist at Leicester and a physicist at Cambridge, Charles Percy Snow (Baron Snow of Leicester) was mid-way through his four years of service as technical director of the Ministry of Labor at the time of this talk.
7. T. Kilburn, "From Cathode Ray Tube to Ferranti Mark I Resurrection," *Resurrection: The Bulletin of the Computer Conservation Society*, vol. 1, no. 2, 1990, pp. 16–20.
8. D.P. Anderson, "Was the Manchester Baby Conceived at Bletchley Park?" *Proc. Electronic Workshops Computing*, British Computer Soc., 2007; http://www.bcs.org/upload/pdf/ewic_tur04_paper3.pdf.
9. D.P. Anderson, "The Contribution of M.H.A. Newman and his Mathematicians to the Creation of the Manchester 'Baby,'" *The British Society for the History of Mathematics Bulletin*, vol. 24, no. 1, 2009, pp. 27–39.
10. D.P. Anderson, "Patrick Blackett: Providing 'White Heat' to the British Computing Revolution" *Communications of the ACM*, Vol. 56 No. 12, Pages 26-28
11. Actually, the process involved developing a series of successively more complex prototype machines.
12. E.M. Dunstan also joined in 1954 and worked on drum storage.
13. S.H. Lavington, *A History of Manchester Computers*, 2nd ed., British Computer Soc., 1998.
14. S.H. Lavington, "The Manchester Mark I and Atlas: A Historical Perspective," *Comm. ACM*, vol. 21, no. 1, 1978, pp. 4–12.
15. D. Morris, "Early Computers at Manchester University," *Computer Resurrection*, vol. 1, no. 4, 1992.
16. T. Kilburn et al., "A System Design Proposal," *IFIP Congress*, vol. 2, 1968, pp. 806–811. See also R.N. Ibbett, "The University of Manchester MU5 Project," *IEEE Annals of the History of Computing*, vol. 21, no. 1, 1999, pp. 24–33.
17. Sir Maurice Wilkes, interview by D.P. Anderson, Feb. 2009. An edited version of this conversation appeared *Communications of the ACM*, Vol. 52 No. 9, Pages 39-42