

Roger Needham : an informal memoir

1935 - 2003

Roger was the only child of Len and Mollie Needham, who met as chemistry students at Birmingham University. He had a happy childhood in Sheffield, going to school there and later in Doncaster. He learnt to read early and was a lifelong and wide-ranging reader, surprising some professional historians, for example.

Roger had a good background for computing. His father was an engineer, familiar with all the problems of making complex systems work (albeit for processing coal); Roger was thoroughly trained in maths (though to the point of surfeit after two years in Cambridge); and he had a year of philosophy (but with the then fashionable post-Wittgensteinian gloss).

In 1956, his last year as an undergraduate at Cambridge, he encountered the Cambridge Language Research Unit, originally a lively discussion group interested in language and translation, subsequently funded to do research on automatic translation. Roger became interested in computing and took the Diploma in Numerical Analysis and Automatic Computing in 1957 (the last year of the EDSAC 1). He worked at the Unit from 1957-1962, also doing research for his PhD. The Unit's projects focussed on the use of a thesaurus, not only for translation but for other information processing tasks like document retrieval, and Roger's research was on automatic classification and its application to (automated) retrieval. Automatic classification was an exciting research area at that time, bringing together people with very different backgrounds and interests but a shared concern with sound general models and computationally viable procedures. Roger applied his theory of clumps to document index terms, lexical data, prehistoric pots, girls' puberty rites, and diseases.

Roger married Karen Spärck Jones in 1958, and they set about building a house, working on site in the mornings and at their respective PhD researches in the afternoons and evenings. In 1961 they bought their first, small and primitive boat.

Roger obtained his PhD in 1961. He had become increasingly interested in computing itself, and joined the staff of the University Mathematical Laboratory (now the Computer Laboratory) as a Senior Assistant in Research in 1963, becoming an Assistant Director of Research in 1964. The Laboratory, under Maurice Wilkes, was engaged with the TITAN project, providing the software for hardware built by ICT/ICL. Roger first worked with David Wheeler on design automation, and then became involved in building the operating system. This was initially a multi-processing system, but subsequently became a multi-access system as well. Much of the work was done in marathon night sessions, or in an office shared with David Hartley and Barry Landy which resembled a prairie dog colony, with heads peering above the mounds of dump printouts. In 1967 he had the idea of storing passwords with a one-way function, and implemented what is now a very common practice.

As a member of staff Roger taught for the Diploma, as he did subsequently for the undergraduate degree course, and began to take PhD students. One student commented on his ability to produce completely well-formed paragraphs without referring to his notes, and also on his ability to mesmerise his audience by walking to and fro while doing so. It has been suggested that his greatest teaching pleasure was in giving the Data Structures and Algorithms course, which he did for many years, taking it as a challenge to interest the students in its necessary but not always sexy content. Roger also began his public service career in

the 1960s as a member of the Science Research Council's Computing Science Committee and, in a different way, as a village Parish Councillor and then South Cambridgeshire District Councillor, continuing with the latter till the 1980s.

In 1966 Roger spent some months as a consultant at the Rand Corporation in Santa Monica, the first in a series of long-standing connections with Californian research centres which he greatly valued, where he had enjoyable collaborations and good friends and colleagues. He was a visiting consultant at Xerox PARC from 1977-84, and at DEC Systems Research Centre from 1984-97. He was involved with the GRAPEVINE project at PARC (on which he later lectured as a case study), and did research with Mike Schroeder on encryption and authentication, published in 1978, and with Mike Burrows and Martin Abadi on a logic of authentication, published in 1989. He was also able, on these visits, to observe the business of running a research centre - how, and also how not, to - at first hand, widening his experience in the 1990s as an advisor for Hitachi's Advanced Research Laboratory.

In the late 1960s Roger began to concentrate more intensively on protection, working with Maurice Wilkes and David Wheeler on the CAP computer, an experimental machine with memory protection based on capabilities implemented in hardware that was constructed in the Laboratory during the early 1970s, and which received a British Computer Society Technical Award in 1977. Roger's research was also affected by organisational changes in the Laboratory. From the days of EDSAC 1, the Laboratory had provided a university-wide computing service on its machines, and the TITAN system was built for this purpose. In 1970 supporting institutional computing was hived off onto IBM machines with a Service division within the Laboratory. Roger remained devoted to the engineering notion that systems should be designed to do useful things for real people, but was able to explore system ideas with less immediate constraints than those imposed by sustaining a large regular service.

The Laboratory's research activity expanded during the 1970s, and Roger was involved in the Cambridge Ring and Fast Ring projects with Maurice Wilkes, David Wheeler and Andy Hopper: the 10 megabit per second Ring and 100 megabits Fast Ring were leading efforts in high-speed local area networking and distributed computing. Roger developed his interest in distributed systems in work with Maurice Wilkes and Andrew Herbert on the Cambridge Model Distributed System, providing an innovative distributed software environment on top of the Ring that prefigured current 'thin client' computing. This line of work was carried further with Ian Leslie and others in the UNIVERSE, and then UNISON, projects in the 1980s on satellite-connected LANs that could support real-time voice and video applications.

Roger had been promoted Reader in Computer Systems in 1973, and when Maurice Wilkes retired in 1980 became Head of Department. He was made Professor in 1981 and elected to the Royal Society in 1985, becoming a Fellow of the Royal Academy of Engineering in 1993. He was awarded the CBE in 2001.

The 1980s were a period of further expansion for the Laboratory. This was the time when government, parsimonious everywhere else, believed in the white heat of information technology and poured in the cash without much restraint. Roger referred to the halcyon days of five new posts, money for research and none of the external interference that, as the audit culture or under the pressure for national wealth creation, has oppressed universities since.

The UNIVERSE and UNISON projects were large collaborations with industrial partners. Roger had had company connections since the 1960s, not only through his Californian consulting, but through lecture courses and discussion clubs (doing what is nowadays referred to as technology transfer), and sometimes as a director. Spinoffs from the Laboratory began in

the 1970s, making significant contributions to the Cambridge Phenomenon. Roger encouraged this when he became Head of Department, welcoming Jack Lang's idea of a Laboratory Supporters Club and becoming one of the 'Godfathers' for Cambridge entrepreneurs.

Roger continued in the 1980s and 90s to be interested in all aspects of computer systems, but was especially concerned with security. He participated for many years in the ACM Symposia on Operating Systems Principles, and was involved, with Ross Anderson, in Cambridge events including a security programme at the Newton Institute and Protocols Workshops. He recently combined his intellectual and (left wing) political interests as a Trustee of the Foundation for Information Policy Research. He also emphasised, in a related spirit, in his 2002 Saul Gorn Lecture at the University of Pennsylvania and Clifford Paterson Lecture at the Royal Society, that doing system security properly is as much about people as about machines.

Though many in the University (and elsewhere) were slow to accept the idea of Computer Science, the Laboratory continued to grow - in staff, undergraduate and graduate students, research projects, and range of interests, during this time. It was also able to take advantage of an anomalous formal position as a department independent of any faculty. The Laboratory has been consistently rated in the top grade in the national university Research Assessment Exercises. Roger was able to feel, when he handed over as Head of Department to Robin Milner in 1996, that the Laboratory was in a flourishing state. He was also happy to step down, as he felt he had been Head for quite long enough.

Roger's public service activities ramified in the 80s and 90s, extending into all kinds of government and other boards and committees. He found some of them fun - the Alvey Committee, for example, had the opportunity to drive a large national computing research programme; some were interesting, like the Research Councils' Individual Merit Promotion Panel; and some were keeping a particular show on the road. Roger felt the obligation to do these things; he also enjoyed learning and deploying the skills required to do them effectively. His most recent challenge was chairing a Royal Society Working Party on intellectual property.

Roger was able to exploit the skills he had developed, and what he had learnt about the University while Head of Department, as Pro Vice-Chancellor from 1996-1998, with a remit on the research side of the University's operations. This had all kinds of interesting side-effects, like chairing Electors to Chairs across the University and so getting snapshots of what's hot in pharmacology, or economic history, or Spanish.

But Roger felt that University life, with ever more auditing and ever more financial pressure spreading like the nastier sort of fungus, was becoming less and less attractive. On retiring, Maurice Wilkes joined DEC, and said that if he had known what fun he could have in industry, he would have done it sooner. Roger never forgot this, and when he was invited in 1997 to set up and run Microsoft's new research laboratory in Cambridge, he seized the opportunity and never looked back. He agreed to do this on the same day that he and Karen viewed and decided to buy a new house (regretfully abandoning their old self built, as too small and too invaded with traffic noise, for a large old barn): as he said: 'I've rebooted my whole life'.

He also said that, contrary to many people's beliefs about companies and universities, he had a lot more freedom to get what he wanted done in his Microsoft Research laboratory than in a university. He had the remit to get good people, and provide them with the context for good work. That implied thinking long-term, and taking risks: he was very happy with the message from Redmond that he wouldn't be doing well enough if none of his projects failed, because that would imply he wasn't pushing the research boat out far enough. At its fifth

anniversary in autumn 2002, the laboratory had about 60 researchers and had moved into its new building in West Cambridge, next to the Computer Laboratory's own new building. The only downside about a company laboratory is not having students, especially bright research students, around all the time. Roger's former students became friends around the world.

Roger loved his work. He also loved sailing. In 1972 Roger and Karen bought an Itchen Ferry Cutter built in 1872, 22 feet long with an 8 foot bowsprit, and for thirty years sailed her round the East Coast, happily but also sometimes rather unnervingly: on one occasion when Roger was singlehanded in a high wind and towing a dinghy, the dinghy just took off from the water and twirled round in the air on the end of its rope.

Roger became a Fellow of Wolfson College, one of the new graduate colleges in Cambridge, in 1966, and always enjoyed its society, just as he enjoyed much else in Cambridge. He and Karen sometimes wondered whether they should take off for greener pastures on the other side of the Atlantic, but found the two body problem too hard to solve (Karen became Reader and subsequently Professor of Computers and Information in the Laboratory, and is a Fellow of the British Academy). They remained colleagues from the time they collaborated and wrote together as students, reading one another's draft papers, working together (most recently in connection with a joint British Academy/Royal Society Discussion Meeting in 1999), and talking about their research.

Late in 2002 Roger was diagnosed with incurable cancer. His colleagues were anxious to celebrate his contributions to computing, and organised a meeting in Cambridge on the 17th of February 2003 - "Roger Needham: 50 and 5" - marking Roger's fifty years in Cambridge and five at Microsoft Research. The meeting, with both technical talks and some more personal items, was attended by colleagues from far and wide. At its conclusion Roger was presented with a volume, 'Computer systems: papers for Roger Needham', with 46 contributions from colleagues in the field, ranging over the whole systems research area. He very much appreciated this evidence of his colleagues' regard, just before he died.

KSJ, 2.3.2003