

Brief Encounters

Brian Randell
University of Newcastle upon Tyne, UK

In compiling these personal reminiscences of Roger, I have - I must confess - mainly relied on my memory rather than consulted my files in any great detail. Given my pretensions to be somewhat of a computer historian, albeit very much part-time, it is somewhat embarrassing to reveal this lack of regard for primary evidence. However, my excuse is that I have no wish to encourage any of you to classify Roger (or me) as historical exhibits.

I cannot recall when Roger's and my paths first crossed - I'm not sure whether they did before I left the UK in 1964 to join IBM Research in Yorktown Heights. Up to this time I'd been working on compiler design for English Electric's DEUCE and then for its KDF9 computers. I had not concerned myself much with computer or operating system design, and had had little contact with Cambridge. So after I arrived at Yorktown Heights I was surprised to learn that one of the people they had contacted to check me out before head-hunting me, so to speak, was Maurice Wilkes.

At Yorktown I made a very deliberate switch from compilers to computer architecture, and this led to me operating systems and so I presume to my first contacts with Roger. I'm not sure now how it happened, but in 1966 I became Editor of the Operating Systems Department of the Communications of the ACM, a post I held for the next seven years. I have tried to find whether I could proudly claim to have accepted any papers by Roger during my seven-year editorial term. As far as I can tell he did not publish in the CACM until after I'd left. Luckily I have no easy way of checking how many of his papers I rejected!

In 1967 I participated in and was editor of the proceedings of the first SOSP, the ACM Symposium on Operating System Principles. The SOSP series, which is one that Roger has had long and extensive involvement with, has in general tried hard to live up to its name and encourage papers that truly do deal with principles - though more recently I gained the impression that it had for a while become somewhat of a mere "Unix Improvements Society" - something I'm sure Roger tried very hard to prevent.

My main memories of the 1967 SOSP symposium, which was held in Gatlinburg, Tennessee, include the fact that at the time Gatlinburg was dry - all you could get within the town limits was something they called "near beer." I presume Roger was in attendance, and remembers this with even more pain than I do, though I believe that his first SOSP paper was at the 3rd Symposium, held in Palo Alto in 1971 - in fact a paper entitled "Handling Difficult Faults in Operating Systems" [1]. This is a paper with a typically high signal-to-noise ratio, covering both principles and practice, whose introductory paragraph is I think worth quoting in full: "It is commonplace to build facilities into operating systems to handle faults which occur in user-level programs. These facilities are often inadequate for their task; some faults or incidents are regarded as so bad that the user cannot be allowed to act on them and this makes it difficult or

impossible to write subsystems which give proper diagnostics in all cases, or which are adequately secure, or which are adequately robust. This paper looks into why there is a need for very complete facilities and why there is a problem providing them, and provides an outline structure which could be used.”

I returned to the UK in 1968, and joined the Computing Laboratory of the University of Newcastle upon Tyne. The Laboratory had been created a dozen years earlier by Ewan Page. Ewan was a Cambridge man, and the Laboratory he set up was I’m sure deliberately patterned after Cambridge’s Computer Laboratory, and similarly combined the roles of academic department and university computing service, a characteristic which we proudly continued to hold and defend for many years, just as Cambridge did. I have often characterized my move from IBM Research to the University of Newcastle upon Tyne with a phrase due to John Buxton, dating from about this time. I cannot resist using an appropriate variant of Buxton’s phrasing to describe Roger’s rather differently directed career move, a few years ago, when what he did was abandon the sordid commercial reality of a university computing laboratory for the ivory towers of industry.

But I’m getting ahead of myself. When I arrived at Newcastle, one of the tasks I took over was the organization of the second in the series of annual Newcastle International Seminars on Computing Science. Roger was one of my choices of speaker for this 2nd Newcastle Seminar. His 1969 talk was on “Failure Recovery” [2], a fact that I must confess that I’d completely forgotten. It would be intriguing to try and determine whether this talk predated the planning we undertook at about this time, following my experiences at the now famous 1968 NATO Software Engineering Conference, that led to our first SERC-sponsored project on dependability, a topic that has been a major feature of my and Newcastle’s research ever since. I have gone back and looked at the report of Roger’s 1969 talk and found that it was about the problems of file system integrity and back-up in the face of unreliable hardware – it includes the nice remark “it is psychologically desirable to take greater care of users’ files than they would themselves” – rather than overall system failure and recovery, so I don’t see any strong link to our early work, which was on techniques for providing continued service despite the presence of residual software faults.

To my surprise, in preparing these remarks, I found that Roger was not present at the 1968 NATO Conference at Garmisch in Bavaria, only at the follow-up conference held a year later in Rome. For various reasons this was a much less effective and influential conference than its predecessor, but Roger made some notable contributions, including one prepared during the conference itself in an intriguing instant collaboration with Joel Aron of IBM Federal Systems Division. My understanding is that Roger and Joel had never met before. Their backgrounds could hardly have been more different – Joel had been heavily involved in the awesome computing system project that supported the Project Apollo series of moon-shots, a project whose scale and style were vastly different from Roger’s work on operating systems in the Cambridge Computer Laboratory. Their styles of speaking were also very different, though in each case very attractive. Joel’s splendid talk on the Project Apollo Ground support system, and each contribution he made to any discussion, always sounded as though he was giving a reading of a carefully structured and punctuated piece of elegant prose. Indeed he did this so clearly that it was child’s play to transcribe a recording of his voice and produce a fully grammatical and properly punctuated text, something that I and the others involved in producing the report

of the conference much appreciated. Roger's style of delivery is on the other hand more notable for its wit and brevity, and thus as entertaining to transcribe as it is to listen to first hand. Yet they somehow found time during a very intense conference to reach a common viewpoint and co-author a paper, albeit a brief one, on "Software Engineering and Computer Science." [3]

At the preceding conference a disparate set of participants, ranging "from the inhabitants of ivory-towered academe to people who were right on the firing line, being involved in the direction of really large scale projects" found "commonality in a widespread belief as to the extent and seriousness of the problems facing the area of human endeavour which has, perhaps somewhat prematurely, been called 'software engineering'." However the report on the Rome Conference, a conference which had a similarly disparate set of participants, comments that "the sense of urgency in the face of common problems was not so apparent as at Garmisch – instead, a lack of communication between different sections of the participants became ... a dominant feature" and explains that "eventually the seriousness of this communication gap, and the realization that it was but a reflection of the situation in the real world, caused the gap itself to become a major topic of discussion." The Aron-Needham paper was a thoughtful contribution to this discussion, and demonstration of the bridgeability of the communication gap, one that I enjoyed re-reading when I prepared these remarks.

Returning to the subject of Roger's contributions to the Newcastle Seminar Series, I should mention that during the thirty-four years of the series, we normally had different speakers each time. Roger is one of the very few speakers who have been invited back not just once but three times – the others being Edsger Dijkstra and Kristen Nygaard. Roger's first reappearance was in 1978. The overall subject of this seminar was Distributed Computing Systems, and Roger talked on "User-Server Distributed Computing" [4]. His two talks provided a very thoughtful analysis of the properties of local area networks such as Ethernet and the Cambridge Ring, and how such networks could be exploited in order to distribute many of the tasks that traditionally were all bundled together into a large monolithic operating system across a set of much simpler specialized servers.

His next appearance was for our 25th seminar, when we used the general title "Computing Science" and deliberately chose our speakers from the by now large set of highly-renowned past speakers. In fact the other speakers at our Silver Jubilee seminar were Edsger Dijkstra, Tony Hoare, Donald Knuth, Butler Lampson, John McCarthy, Kristen Nygaard, and Michael Rabin. Almost all of the Jubilee speakers fully lived up to their reputations and gave excellent talks - Roger certainly did, with talks on "Communication System Development," and on "Reasoning about Cryptographic Protocols" [5]. This latter was of course largely based on his very influential and much-cited work with Michael Burrows and Martin Abadi on the BAN logic, the notation they designed for use in analyzing and verifying authentication protocols [6]. To complete the list, I should mention that Roger was a speaker at the last in the Seminar Series, in September 2001 [7] – this was a sort of benefit match for me, since it was on Dependability, and marked my (so-called) retirement. It was the first seminar in over thirty years that I had not organised – my colleague and professorial successor, Cliff Jones, was in charge – and thus for the first time ever in the series I found myself having to lecture. But who better could Cliff have chosen to speak on Security than Roger?

But again I'm getting ahead of myself. Following my return to the UK, I had many opportunities to meet up with Roger. For example we served together for what seemed like many years on a whole succession of DTI advisory committees. Though at times this was an enjoyable experience – because Roger has an inimitable way of speaking to and dealing with recalcitrant civil servants, one that I find much more entertaining than they do – it was also, we both agree, an immensely frustrating experience. Though we were not so naïve as to assume that all our advice would be heeded, it is clear in retrospect that we were almost entirely wasting our time. No wonder that, as I've since learned, the Department is referred to by some as the Department of Timidity and Inaction.

The one exception, the one really worthwhile experience I had with the DTI, was again one I shared with Roger. This was on a 1981 DTI mission to Japan. There is a fairly full description of this mission in the book "Alvey: Britain's Strategic Computing Initiative" by Brian Oakley and Kenneth Owen [8]. One entertaining (and all too true) passage is the following: "On arrival in Tokyo team members were fascinated to discover a completely alien culture, with strange customs, exotic behaviour, and quaint patterns of speech. And that, they recall, was just the British Embassy."

We (namely Roger, Alan Fox of RSRE Malvern as it then was called, Charles Read of the Inter-Bank Research Organization, Reay Atkinson, an uncharacteristically splendid DTI civil servant, and I) had been sent out to Japan by a Minister who had come back from a visit completely entranced and highly flattered by the Japanese government's invitation to the UK to participate in the Fifth Generation Computer Project that they were planning. The Alvey book attributes to me the following subsequent assessment of the Japanese plans: "The Japanese conference presentations were an amazingly well-orchestrated series of vague accounts of various parts of an ambitious and wide-ranging plan . . . Everybody made respectful references, at least, to logic programming and knowledge engineering, and some of them obviously believed, and perhaps even understood, what they were saying. It came over to me as a very skilful plan which filled MITI's wish for a very ambitious goal that sounded very plausible and which could be presented to a layman in such a way as to seem socially beneficial."

However, though attributed to me, I think I detect the hand of Roger in that text – if not, it is the effect on me of an extended period of close proximity to him. The team rapidly came to the conclusion that we wished to dissuade the government from setting up a general scheme of UK-Japanese collaboration, since it was clear to us that the main beneficiaries would be Japanese industry. (We had discovered from visits we paid to various Japanese computer companies that they all were much better informed on the latest UK academic computer science research than any UK company.) Instead, we argued, what was first needed was some effective means of encouraging collaboration between UK academics and industry on a large-scale programme of Information Technology R&D, a programme which should not be so narrowly focussed on logic programming and knowledge engineering as the Japanese 5G plans.

After our return to the UK we were thus both heavily involved in the scheming that led to the Alvey Programme, but that is a whole story by itself, and one that has already been well-documented. However, it is important to point out how central was the role that Roger played in the setting up of the Alvey Programme, not least as the sole academic allowed to join the committee of senior industrialists and civil servants, chaired by Sir John Alvey, whose report directly led to the creation of this ground-breaking programme.

The Mathematical Sciences Sub-Committee of the late lamented University Grants Committee was another arena that has provided me with happy memories of encounters with Roger. There was for example the sub-committee visit to a particular university (fortunately I cannot remember which one) when, during the obligatory tour of the CS department aimed at gaining our support for additional accommodation, Roger became aware that the faces in the various laboratories we inspected were becoming familiar. This was because a crowd of students was being rushed round back corridors of the department to reappear in front of us repeatedly, rather like the chorus during the Grand March in an under-staffed performance of the opera Aida. Such visits also normally involved meetings with the local Vice-Chancellor – at which Roger demonstrated a skill in the sometimes rather delicate discussions which I’m sure proved very useful when he became a Pro-Vice-Chancellor himself at Cambridge some years later.

During this time, the UGC Mathematical Sciences Sub-Committee, under the able leadership of Prof. Douglas Jones, a canny Scot if ever there was one, was involved not just in a regular programme of such visits, but also in considering submissions from just about all the UK computer science departments to a whole succession of funding initiatives. We were thus able to gain detailed knowledge of the then fifty or so departments, as well as to achieve a number of significant resource enhancements for UK computer science. As a result, when we learnt that we were going to have to perform what turned out to be the first of the dreaded Research Assessment Exercises, Roger and I independently drew up, while waiting for our trains at Kings Cross Station, virtually identical and remarkably accurate predictions of the gradings that later resulted from the subsequent formal consideration of the detailed RAE submissions by the overall panel. To paraphrase a comment once made about Algol 60, this first RAE exercise was in my opinion, in regard both to the way it was carried out, and the degree of acceptance of the results by the UK computer science community, an improvement over all its successors.

Enough of committees and bureaucracy – let me end with a few further remarks on research. Roger’s and my research trajectories have diverged somewhat over the last thirty years. He has concentrated largely on security issues – to great effect – whereas I’ve worked on fault tolerance, as applied to reliability and availability, though I have on occasion had fun investigating potential links between fault tolerance and security. However, in the early days Roger was equally interested in what was essentially fault tolerance even if he didn’t use this term – I recall an early aphorism of his to the effect that operating systems should be designed and implemented via incremental additions to a very robust dump and restart system. I have enjoyed looking back at a number of his early papers – even if I now have some concerns as to whether I paid them as much attention at the time as they evidently deserved. Thus I can sympathize very much with the comment that Roger made in an interview in 2001: “Although for most of my career I was a practical builder of systems, the things I’m best known for are [two papers on authentication], both of a theoretical nature and both done when I was on sabbatical leave. So you can work away on a complicated system for seven years, and nobody remembers that.” [9] (Incidentally, during my little investigation of Roger’s early papers, I was startled to find that he had published one in 1964 entitled “Exploitation of Redundancy in Programmes” [10] – however, this turned out to be concerned with instruction set representation, and the issues discussed were instruction storage efficiency and processor performance, not dependability!)

Regarding security, Roger's expertise regarding cryptography of course far exceeds mine – in fact I'm sure my evident lack of knowledge of, or interest in, cryptography was of considerable benefit when I was seeking official permission to investigate Britain's highly classified wartime code-breaking machines, in particular the Colossus. But we both share a degree of scepticism about the subject of cryptography, and I very much like the comment often attributed to Roger, including by Butler Lampson, though I gather Roger claims it was Butler who first said that: "Anybody who asserts that a problem is readily solved by encryption, understands neither encryption nor the problem." [11]

To be with Roger is to enjoy, and benefit from, a whole succession of such wise and pithy remarks – it is thus a great pleasure to place on record how highly I value all the opportunities I've had of encounters with Roger from time to time over the years, and all the enjoyment and benefit I've thus gained.

References

1. NEEDHAM, R.M., 'Handling difficult faults in operating systems.' in *Proc 3rd Symp. on operating system principles (Operating Systems Review, vol. 6, no. 1, 2, June 1972)*, pp. 55-57, Palo Alto, CA, 1971.
2. NEEDHAM, R.M., 'Failure recovery,' in *On the teaching of the design of large software systems (Proc. Joint IBM/University of Newcastle upon Tyne Seminar, 8-12 September 1969)*, ed. N. S. M. Cox, pp. 38-43, University of Newcastle upon Tyne, 1970.
3. NEEDHAM, R.M. AND J.D. ARON, 'Software engineering and computer science,' in **Software engineering techniques: report on a conference sponsored by the NATO science committee**, Rome, Italy, 27-31 Oct. 1969, ed. J. N. Buxton and B. Randell, pp. 113-114, Brussels, NATO Science Committee, 1970.
4. NEEDHAM, R.M., 'User-server distributed computing,' in *Distributed computing systems*, ed. B. Shaw, pp. 71-78, Computing Laboratory, University of Newcastle upon Tyne, 1987.
5. Needham, R.M., 'Reasoning about cryptographic protocols,' in *ed. B. Randell*, Department of Computing Science, University of Newcastle upon Tyne, 1992.
6. BURROWS, M. ABADI, M. AND NEEDHAM, R.M., 'A logic of authentication,' *ACM Transactions on Computer Systems*, vol. 8, no. 1, pp. 18-36, 1990.
7. NEEDHAM, R.M., 'Security,' in *Dependability*, ed. C. B. Jones, Department of Computing Science, University of Newcastle upon Tyne, 2001.
8. OAKLEY, B. AND OWEN, K., **Alvey: Britain's strategic computing initiative**, Cambridge, MA, MIT Press, 1989, 337p.
9. OMITOLA, T., 'ACM Fellow Profile: Roger Needham,' *Software Engineering Notes*, vol. 26, no. 1, 2001.
10. NEEDHAM, R.M., 'Exploitation of redundancy in programmes,' in *Proc. Conf. on the impact of users' needs on the design of data processing systems*, pp. 6-7, UK, IEE, 1964.
11. NEEDHAM, R.M., 'Address,' in *The Marshall symposium - the information revolution in midstream: an angl- american perspective*, Rackham School of Graduate Studies, University of Michigan, 1998. (<http://www.si.umich.edu/marshall/docs/p201.htm>)