

REMINISCENCES ON THE HISTORY OF TIME SHARING

John McCarthy, Stanford University

1983 Winter or Spring

I remember thinking about time-sharing about the time of my first contact with computers and being surprised that this wasn't the goal of IBM and all the other manufacturers and users of computers. This might have been around 1955.

By time-sharing, I meant an operating system that permits each user of a computer to behave as though he were in sole control of a computer, not necessarily identical with the machine on which the operating system is running. Christopher Strachey may well have been correct in saying in his letter to Donald Knuth that the term was already in use for time-sharing among programs written to run together. This idea had already been used in the SAGE system. I don't know how this kind of time-sharing was implemented in SAGE. Did each program have to be sure to return to an input polling program or were there interrupts? Who invented interrupts anyway? I thought of them, but I don't believe I mentioned the idea to anyone before I heard of them from other sources.

My first attempts to do something about time-sharing was in the Fall of 1957 when I came to the M.I.T. Computation Center on a Sloan Foundation fellowship from Dartmouth College. It was immediately clear to me that the time-sharing the IBM 704 would require some kind of interrupt system. I was very shy of proposing hardware modifications, especially as I didn't understand electronics well enough to read the logic diagrams. Therefore, I proposed the minimal hardware modification I could think of. This involved installing a relay so that the 704 could be put into trapping mode by an external signal. It was also proposed to connect the sense switches on the console in parallel with relays that could be operated by a Flexowriter (a kind of teletype based on an IBM typewriter).

When the machine went into trapping mode, an interrupt to a fixed location would occur the next time the machine attempted to execute a jump instruction (then called a transfer). The interrupt would occur when the Flexowriter had set up a character in a relay buffer. The interrupt program would then read the character from the sense switches into a buffer, test whether the buffer was full, and if not return to the interrupted program. If the buffer was full, the program would store the current program on the drum and read in a program to deal with the buffer.

It was agreed (I think I talked to Dean Arden only.) to install the equipment, and I believe that permission was obtained from IBM to modify the computer. The connector to be installed in the computer was obtained.

However, at this time we heard about the "real time package" for the IBM 704. This RPQ (request for price quotation was IBM jargon for a modification to the computer whose price wasn't guaranteed), which rented for \$2,500 per month had been developed at the request of Boeing for the purpose of allowing the 704 to accept information from a wind tunnel. Some element of ordinary time-sharing would have been involved, but we did not seek contact with Boeing. Anyway it was agreed that the real time package, which involved the possibility of interrupting after any instruction, would be much better than merely putting the machine in trapping mode. Therefore we undertook to beg IBM for the real time package. IBM's initial reaction was favorable, but nevertheless it took a long time to get the real time package - perhaps a year, perhaps two.

It was then agreed that someone, perhaps Arnold Siegel, would design the hardware to connect one Flexowriter to the computer, and later an installation with three would be designed. Siegel designed and build the equipment, the operating system was suitably modified (I don't remember by whom), and demonstration of on-line LISP was held for a meeting of the M.I.T. Industrial Affiliates. This demonstration, which I planned and carried out, had the audience in a fourth floor lecture room and me in the computer room and a rented closed circuit TV system. Steve Russell, who worked for me, organized the practical details including a rehearsal. This demonstration was called time-stealing, and was regarded as a mere prelude to proper time-sharing. It involved a fixed program in the bottom of memory that collected characters from the Flexowriter in a buffer while an ordinary batch job was running. It was only after each job was run that a job that would deal with the characters typed in would be read in from the drum. This job would do what it could until more input was wanted and would then let the operating system go back to the batch stream. This worked for the demonstration, because at certain hours, the M.I.T. Computation Center operated at certain hours a batch stream with a time limit of one minute on any job.

Around the time of this demonstration, Herbert Teager came to M.I.T. as an assistant professor of Electrical Engineering and expressed interest in the time-sharing project. Some of the ideas of time-sharing overlapped some ideas he had had while on his previous job, but I don't remember what they were. Philip Morse, the Director of the Computation Center, asked me if I was agreeable to turning over the time-sharing project to Teager, since artificial intelligence was my main interest. I agreed to this, and Teager undertook to design the three Flexowriter system. I'm not sure it was ever completed. There was a proposal for support for time-sharing submitted to NSF and money was obtained. I don't remember whether this preceded Teager, and I don't remember what part I had in preparing it or whether he did it after he came. This should be an important document, because it will contain that year's conception of and rationale for time-sharing.

Besides that, IBM was persuaded to make substantial modifications to the IBM 7090 to be installed at the M.I.T. Computation Center. These included memory protection and relocation and an additional 32,768 words of memory for the time-sharing system. Teager was the main specifier of these modifications. I remember my surprise when IBM agreed to his proposals. I had supposed that relocation and memory protection would greatly slow the addressing of the computer, but this turned out not to be the case.

Teager's plans for time-sharing were ambitious and, it seemed to many of us, vague. Therefore, Corbato undertook an "interim" solution using some of the support that had been obtained from NSF for time-sharing work. This system was demonstrated some time in 1962, but it wasn't put into regular operation. That wasn't really possible until ARPA support for Project MAC permitted buying a separate IBM 7090.

Around 1960 I began to consult at BBN on artificial intelligence and explained my ideas about time-sharing to Ed Fredkin and J. C. R. Licklider. Fredkin, to my surprise, proposed that time-sharing was feasible on the PDP-1 computer. This was D.E.C.'s first computer, and BBN had the prototype. Fredkin designed the architecture of an interrupt system and designed a control system for the drum to permit it to be used in a very efficient swapping mode. He convinced Ben Gurley, the chief engineer for D.E.C. to build this equipment. It was planned to ask NIH for support, because of potential medical applications of time-sharing computers, but before the proposal could even be written, Fredkin left BBN. I took technical charge of the project as a one-day-a-week consultant, and Sheldon Boilen was hired to do the programming. I redesigned the memory extension system proposed by D.E.C. and persuaded them to build the modified system instead of the two systems they were offering, but fortunately hadn't built. I also supervised Boilen.

Shortly after this project was undertaken, D.E.C. decided to give a PDP-1 to the M.I.T. Electrical Engineering Department. Under the leadership of Jack Dennis, this computer was installed in the same room as the TX-0 experimental transistorized computer that had been retired from Lincoln Laboratory when TX-2 was built. Dennis and his students undertook to make a time-sharing system for it. The equipment was similar, but they were given less memory than the BBN project had. There wasn't much collaboration.

My recollection is that the BBN project was finished first in the summer of 1962, but perhaps Corbato remembers earlier demonstrations of CTSS. I left for Stanford in the Fall of 1962, and I hadn't seen CTSS, and I believe I hadn't seen Dennis's system operate either. BBN didn't operate the first system and didn't even fix the bugs. They had few computer users and were content to continue the system whereby users signed up for the whole computer. They did undertake a much larger follow-on project involving a time-shared PDP-1 that was installed in Massachusetts General Hospital, where it was not a success. The computer was inadequate, there were hardware and software bugs, and there was a lack of application programs, but mainly the project was premature.

At the same time that CTSS, the BBN system, and the EE Department systems were being developed, M.I.T. had started to plan for a next generation computer system. The management of M.I.T. evidently started this as an ordinary university planning exercise and appointed a high level committee consisting of Philip Morse, Albert Hill and Robert Fano to supervise the effort. However, the actual computer scientists were persuaded that a revolution in the way computers were used - to time-sharing - was called for. The lower level committee was chaired by Teager, but after his ideas clashed with those of everyone else, the committee was reconstituted with me as chairman. The disagreement centered around how ambitious to be and whether to go for an interim solution. Teager wanted to be very ambitious, but the rest of us thought his ideas were vague, and he wanted M.I.T. to acquire an IBM 7030 (Stretch) computer as an interim solution. As it turned out, acquiring a Stretch would have been a good idea.

Our second report to M.I.T. proposed that M.I.T. send out a request for proposals to computer manufacturers. On the basis of the responses, we would then ask the Government for the money. The RFP was written, but M.I.T. stalled perhaps for two reasons. The first reason was that our initial cost estimates were very large for reasons of conservatism. Secondly, IBM asked M.I.T. to wait saying that they would make a proposal to meet M.I.T.'s needs at little or no cost. Unfortunately, the 360 design took longer than IBM management expected, and along about that time, relations between M.I.T. and IBM became very strained because of the patent lawsuit about the invention of magnetic core memory.

As part of the stall, President Stratton proposed a new study with a more thorough market survey to establish the demand for time-sharing among M.I.T. computer users. I regarded this as analogous to trying to establish the need for steam shovels by market surveys among ditch diggers and didn't want to do it. About this time George Forsythe invited me to come back to Stanford with the intention of building a Computer Science Department, and I was happy to return to California.

In all this, there wasn't much publication. I wrote a memo to Morse dated January 1, 1959 proposing that we time-share our expected "transistorized IBM 709". It has been suggested that the date was in error and should have been 1960. I don't remember now, but I believe that if the memo had been written at the end of 1959, it would have referred to the 7090, because that name was by then current. In that memo I said the idea of time-sharing wasn't especially new. I don't know why I said that, except that I didn't want to bother to distinguish it from what was done in the SAGE system with which I wasn't very familiar.

Most of my argumentation for time-sharing was oral, and when I complained about Fano and Corbato crediting Strachey with time-sharing in their 1966 Scientific American article, Corbato was surprised to find my 1959 memo in the files. Their correction in Scientific American was incorrect, because they supposed that Strachey and I had developed the idea independently, whereas giving each user continuous access to the machine wasn't Strachey's idea at all. In fact, he didn't even like the idea when he heard about it.

Teager and I prepared a joint abstract for an ACM meeting shortly after he arrived, and I gave a lecture in an M.I.T. series called Management and the Computer of the Future. In this lecture I referred to Strachey's paper "Time-sharing of large fast computers" given at the 1959 IFIP Congress in Paris. I had read the paper carelessly, and supposed he meant the same thing as I did. As he subsequently pointed out, he meant something quite different that did not involve a large number of users, each behaving as though he had a machine to himself. As I recall, he mainly referred to fixed programs, some of which were compute bound and some input-output bound. He did mention debugging as one of the time-shared activities, but I believe his concept involved one person debugging while the other jobs were of the conventional sort.

My 1959 memo advertised that users generally would get the advantage of on-line debugging. However, it said nothing about how many terminals would be required and where they would be located. I believe I imagined them to be numerous and in the users' offices, but I cannot be sure. Referring to an "exchange" suggests that I had in mind many terminals. I cannot now imagine what the effect was on the reader of my failure to be explicit about this point. I'm afraid I was trying to minimize the difficulty of the project.

The major technical error of my 1959 ideas was an underestimation of the computer capacity required for time-sharing. I still don't understand where all the computer time goes in time-sharing installations, and neither does anyone else.

Besides M.I.T.'s NSF proposal, there ought to be some letters to IBM and perhaps some IBM internal documents about the proposal, since they put more than a million dollars worth of equipment into it. Gordon Bell discusses D.E.C.'s taking up time-sharing in Bell and Newell book, but I don't recall that they discuss Ben Gurley's role. Fredkin and perhaps Alan Kotok would know about that.

After I came to Stanford in 1962, I organized another PDP-1 time-sharing project. This was the first time-sharing system based on display terminals. It was used until 1969 or 1970 for Suppes's work on computer aided instruction. [1994 note: Then it was donated to the Indian Institute of Technology at Kanpur, where it was used for about 10 years.]

Appendix:

Don Knuth, who was curious about who had done what, wrote to Christopher Strachey and got the following reply.

OXFORD UNIVERSITY COMPUTING LABORATORY 45 Banbury Road
PROGRAMMING RESEARCH GROUP Oxford OX2 6PE
1st May 1974

Professor D. E. Knuth
Stanford University
Computer Science Department
Stanford, California 94305
U.S.A.

Dear Don:
The paper I wrote called 'Time Sharing in Large Fast Computers' was read at the first (pre IFIP) conference at Paris in 1960. It was mainly about multi-programming (to avoid waiting for peripherals) although it did envisage this going on at the same time as a programmer was debugging his program at a console. I did not envisage the sort of console system which is now so confusingly called time sharing. I still think my use of the term is the more natural.

I am afraid I am so rushed at the moment, being virtually alone in the PRG and having just moved house, that I have no time to look up any old notes I may have. I hope to be able to do so while settling in and if I find anything of interest I will let you know.

Don't place too much reliance on Halsbury's accuracy. He tends to rely on memory and get the details wrong. But he was certainly right to say that in 1960 'time sharing' as a phrase was such in the air. It was, however, generally used in my sense rather than in John McCarthy's sense of a CTSS-like object.

Best wishes,
Yours sincerely,
C. Strachey
Professor of Computation
University of Oxford

/@steam.stanford.edu:/u/ftp/timesharing.tex: begun 1994 Dec 29, latexed

JanFebMarAprMayJun JulAugSepOctNovDec at :< 10 0