M.I.T. LABORATORY FOR COMPUTER SCIENCE
Computer Systems Research

September 5, 1980
Request for Comments No. 197

TRIP REPORT, AUGUST 15 - SEPTEMBER 1, 1980

by Jerome H. Saltzer

My itinerary was: IBM Zurich Research Laboratory, 2 days; Cambridge University Computer Laboratory, 2 days; British Post Office Research group, Cambridge, 1 day; Thomson CSF, Paris, 1 day (cancelled); IFIP WG 6.4 Zurich Workshop on local area computer networks, 3 days.

1. IBM Zurich Research Laboratory. This lab has undergone a noticeable change in the last three years, as the new manager of computer science research (Martin Reiser) has focused the staff on projects with less emphasis on communication technology and more emphasis on data communication systems. A prototype buffered star switch that operates at an internal speed of 150 mb/sec. was reported in a paper at the local net workshop (described later) and a 4 mb/sec. token ring network with synchronous timing set by some single station is under development. This ring design is quite interesting, as it draws ideas from both the LCS ring and the Cambridge ring, yet is being developed in an industrial environment. Some rather sophisticated engineering analyses were undertaken to support the ring design; a paper to be delivered in November describes part of this work. A transport system interface quite similar in spirit to our own and somewhat more independent of the underlying technology is being designed; recognition of the need for a standard interface at this position in a local network was also expressed at the local net workshop.

2. University of Cambridge Computer Laboratory. My host there was Professor Maurice Wilkes, who was in his last week at the Computer Laboratory before retiring from there; he is taking an assignment at the Digital Equipment Corporation in Maynard, Mass., starting September 1, 1980. I was taken on a perambulation of the Cambridge ring, which operates synchronously at 8 mb/sec. and had 31 nodes attached to it. The National Research Council has encouraged other universities to put the ring into production use; several are in the process of doing so, led by the

University College London (Peter Kirstein) and the University of Kent (Brian Reed). UCL has converted the ring interface to a printed circuit board to lower the station cost, and Cambridge (Andy Hoppar) is about to receive a special LSI chip (being constructed from a gate array), also to lower station cost.

Other places installing (or selling) Cambridge rings are Linotype Paul, Logica, Raycal, and 12 other British universities.

The ring uses unshielded twisted pair at Cambridge, and apparently generates some noise to other nearby equipment. At the University of Kent, they have installed shielded twisted pair, and the bit error rate dropped from 1 in 5×10^{11} (e.g., one bit error every day) to unmeasureable, over 100-meter links, and no cross talk problem has yet been found. Jitter contributed by a single repeater in the Cambridge system is below 1 ns out of the 100 ns bit frame when idling (all zeros for data) and rises to 6 ns for random data. (The 4-wire system is limited in distance by skew among the drivers and wires, and it was not clear whether the jitter is absolute or interwire, since it was claimed that the jitter is damped out at every receiving station.) All these observations confirm our own measurements and are consistent with our design criteria for the 8 Mbit/sec. LCS ring net.

The primary project at Cambridge is a distributed computing experiment, under guidance of Roger Needham, built on the Cambridge ring, with basic architecture consisting of a bank of small or medium-sized processors that can be assigned dynamically to support login sessions from terminals. This approach is intermediate between that of the traditional time-sharing system and that of a fully dedicated personal computer, and is based on an assumption that the cost of a processor and enough supporting memory will remain too high for desktop computers for a while.

Storage is accomplished by a file server, (project of Jeremy Dion) which has a traditional read/write interface with transaction/commit semantics added. The file server is operating, and working papers describing its design are just now being written. This file server is one that uses garbage collection techniques as an integral part of the design.

An authentication server based in part on the Needham and Schroeder protocols, extended with a capability view, is being designed by a graduate student, Greg Girling.

3. British Post Office. The BPO, which is responsible for telephone and data communication and the Viewdata service, has a small laboratory in Cambridge; people from the main laboratory at Ipswich and from London came to Cambridge to meet with me. This session ranged over a wide variety of topics in which we established a large number of mutual interests. Much of our discussion centered about Viewdata, which now operates from 20 centers. This application is interesting because although it is distributed, it is specialized enough that update can be accomplished by simple algorithms. One viewdata center is an update collection machine; following an update the new information is broadcast to the 20 service centers (which are otherwise read-only) typically within a few minutes. It appears that the primary roadblock to success of Viewdata is that the British television manufacturers have so far resisted making any investment in the concept; their current prices (typically a $700 premium on top-of-the-line $700 sets only) tend to be discouraging to widespread use.

The British Telephone system still is almost entirely step-by-step; they have managed to omit both the crossbar and electronic switch generation of telephone equipment. The primary question now is whether or not to leap frog feet first into all-digital systems merging voice and data. The opportunity is being eyed warily by those nervous that not enough local experience has been gained to make such a leap. The next ten years of activity should be interesting to watch. The BPO has, with its data customers, helped develop a recommended end-to-end transport protocol that allows attachment of local nets to the X.25 public data network, using source routing, among other strategies. (Source routing is suggested on the basis that it appears hopeless to expect local network suppliers to agree on a addressing standard.) A BPO publication describing this suggested protocol, and known as the "Yellow book", is just now being circulated. It appears that industrial and university users of data communication in Britian have contributed to this protocol and intend to implement it quite widely. It fits into the ISO architecture recommendation at level 4, and is quite different from one being proposed by ECMA.

4. My planned visit to Thomson-CSF, a major telephone company supplier in Paris, was cancelled the day before my arrival after it was established that too many of the people I should meet were on vacation. Although some confusion on dates aggravated the problem, it appears that it is not a good idea to attempt to schedule laboratory visits to Paris in August, since almost every Parisian takes much of the month for vacation. In addition, most of the restaurants in Guide Michelin are closed.

5. IFIP Working Group 6.4 and IBM Zurich Laboratory local area computer network workshop. This workshop attracted workers from every important local network project that I know of, (and a few I didn't know of,) so it was an excellent opportunity to get personally updated, and to let others know what our laboratory is doing. There were 93 attendees from 10 countries, including 33 Americans. I learned lots of random things.

   - CMU has taken delivery on one PERQ. It came with no network, and it appears that the manufacturer, Three Rivers, may be giving the network a lower priority than the machine itself. Some people at CMU are wondering about an alternative local network.

   - A company named Ungermann and Bass is installing 4 mb/sec. Ethernets with RS232 interfaces. An engineer named John Davidson was going after the workshop to Lausanne, Switzerland, to install one.

   - The NBS Ethernet runs at 1 mb/sec., and seems to cost $3K per station, comparable to other such designs and our own higher speed ring cost. They report many problems with passive components (e.g., connectors) and 70% of their power supplies failed in the first year. Installation costs (not pulling the cable, but rather getting the node ready) are running much higher than anticipated. This site reports a lightning strike to its local net, along with Xerox PARC and Cambridge.

   - The University of Kent reports that in 5 months of operating a copy of the Cambridge ring they have had only a couple of quickly isolated failures. They believe that a ring can be very reliable, since theirs is.

- Xerox has apparently helped a consortium of cable manufacturers get UL approval for Teflon-filled coaxial cables for installation in air plenums. The first "Ethernet" cables are expected to appear in this fall's catalog of at least one cable manufacturer.
- There was a repeated, high interest in a "technology independent interface" from people who want to allow either an Ethernet or a ring to be installed at their site. There doesn't seem to be any suitable available standard; candidates such as RS449 are not clearly applicable at speeds as high as 10 mb/sec. The proposed Xerox/Dec/Intel interface is not technology independent-- you can tell you are talking to an Ethernet.

There was a special session devoted mostly to the Xerox/Dec/Intel proposed Ethernet standard, at which a few things were said publicly for the first time. The 10 mb/sec. network is allowed to cover a maximum end-to-end distance of 2.5Km in 500-meter segments connected with repeaters. From 100 to 1000 stations can be attached, depending on their individual noise contribution specification; the network can be branching. (These specifications sound to me to be very hard to meet all at once.) Some reasons were given for not using HDLC/ADCCP logical structure: interactions between the logical and physical layer (e.g., carrier presence/absence eliminates the need for framing patterns and consequently bit stuffing) a feeling that 48-bit addresses should be uniformly used, an opinion that low-level acknowledgements and priority are complicated and unnecessary at this level. A detailed set of Ethernet specifications (being written by David Redell) is scheduled to be available by the end of September. The Intel representative described a planned chip that provides packet buffering, DMA to a Multibus, buffer and command chaining, scatter read/write, programmable interrupts, and test modes. The chip is in HMOS technology, targeted as a 1982 product, and uses external clock recovery. It was claimed (privately) that the only reason for limiting the Ethernet speed to 10 mb/sec. was the ability of building an HMOS chip that runs fast enough.

The 20 papers and summaries of the workshop sessions will be published as a book, as I recall, by North-Holland.