****

Report of the EF-57 study tour covering visits to major semiconductor companies, data processing companies & concefences in USA.

****

NOTE: This report contains information obtained under non-disclosure agreements and confidentiallity agreements. **The content is therefore solely intended for the people involved in the EF-57 project and must not be disclosed to any third part without the express written consent of the author.**

Tue Bertelsen

SDC .

received by: _____

date: _____

Conclusion
****

This report covers a visit to the United States in October and November 1981. During this tour, several conferences on distributed data processing and local computer networks were attended. Visits were made to several of the major semiconductor companies in "Silicon Valley" to study the progress of semiconductor VLSI technology and its impact on our future data processing systems. Visits were also made to some smaller companies representing technological know-how in other fields than pure semiconductor manufacturing, and to some mainframe equipment manufacturers.

The basic outcome of the tour was an impressive confirmation of the results and conclusions hitherto obtained in this project. Also, SDC and its microprocessor development group was made known in Silicon Valley, and maybe the far most important personal outcome was the fact that we can discuss and contribute in significant areas of the American computer development.

The tour covered a period of one and a half month, which turned out to be too short a period for this type of visit. Nevertheless, the outcome should prove worthwhile for both this project and SDC´s future developments as well.

Visit at Beehive International,
4910 Amelia Earhart Drive, Salt Lake City.
****

This visit had been arranged through ScMetric in Denmark. It was the very first visit on the tour, and as my plane from Copenhagen was delayed 4 hours, I missed my connecting flight to Salt Lake City, and had to stay in Seattle overnigth. My meeting at Beehive was scheduled for the following friday morning. Friday morning I succeeded in getting contact with Beehive from Seattle, and we postponed the meeting to the late afternoon. When I finally came to Salt Lake City, I was picked up by an employee from Beehive, so I had no trouble in finding the place. Later, I found out that they actually had been in the airport the previous evening to meet me, and when I didn't arrive, they had driven to my hotel to leave a message there. I never found out, whether this friendly nature was due to their mormon religion or just extraordinary concern for their customers.

Originally, I should have met mr. John Bodell, Director of Sales, but due to the change of schedule, he was not present in the afternoon, so I talked with

Bruce T. Stonely Product Support

plus a couple of other employees, whom I didn't get the names of. Actually, the company seemed to be almost closed when I arrived, and I found out, that most of the employees at Beehive had a four day working week, with 10 hours of work each day, so they could take the friday off.

The purpose of the meeting was to discuss the trends in VDU design and the impact of technology on the architecture and facilities of VDU terminals. Beehive is one of the leading companies in providing very cost/effective smart terminals with built-in features suited for business applications.

I gave a presentation of SDC, the current online system and our migration to the TP2 system. I also explained our current use of Beehive terminals in connection with our microprocessor developments, and explained some of our thoughts behind our requirements for the VDU terminals

delivered by Olivetti.

I was then presented the current and some of the future Beehive products. One of the new products, the Microbee 4400, was actually shown, and even if its physical design seemed a bit clumsy compared to the current Microbee series and to European taste, it contained a number of interesting features.

The design included a user-programmable set-up of the configuration of the terminal. The set-up was menu-driven, where the user was prompted for the various details in the set-up, like syncronous/asynchronous operation, half/full duplex mode, number of stop bits, ASCII/EBCDIC mode, current loop or RS-232 mode, function key assignment etc. The set-up could be stored in non-volatile memory, so it could automatically be restored at power-on.

Also, the terminal contained a number of internal self-diagnostic programs, which was activated during the power-on sequence, and which could be initiated either through the keyboard or through a remote host via the communication link.

We then had a long discussion on the architecture of future terminals based on VLSI technology. Beehive stated, that VLSI technology would become an integral part of terminals, both in order to cope with the competition from other manufacturers, to lower actual maintenance cost by providing a simpler design based on fewer components and more use of self-test capabilities, and to provide the user with the possibility of a more costumer-tailored device due to the programmability.

Beehive was virtually shocked to hear about our TP2-system and the way the VDU terminals was integrated into the terminal computer, as they believed the concept of smart or intelligent terminals supporting a standard interface to be the only survivable method in the future. One of the indicators of this trend was IBM's move towards the support of RS-232 compatible ASCII terminals.

Beehive recommended us to follow the activities within the local network standards efforts very closely. They were themselves designing an Ethernet compatible terminal, but

would not give any indications of production schedules. They felt however, that most companies would come out with Ethernet compatible devices within very few years.

Concerning VLSI-based architectures, Beehive informed me that they had established a joint-venture with one of the semiconductor companies, where Beehive had done the functional specifications on a set of three VLSI chips, constituting the entire logic of a smart VDU terminal. They felt, that it was more advantageous to provide a terminal with programmability of functional behaviour rather than try to do a specialized design of e.g. a graphics terminal. Once having designed a general purpose chip set, it would only be a matter of programming and memory capability to provide a given terminal with graphics or colors.

Internally, they had done some computer design, just be sure that they were able to design other things than terminals, but they had no intentions of producing a computer themselves.

We discussed, how the human interface could be improved by addtional features in the terminals, and Beehive stated that they actually had made no kind of instructions on how to use VDU features like reverse video, half-tone fields etc. optimally. They felt, that this was individually from user to user, but that it was important, that the terminals contained such features from the birth.

We talked about the ergonomic issues when designing terminals, and I was told that the Americans didn't like the brown color of the current microbee series, but preferred the old light and dark grey colors. They were interested in the photographs of the TP2 terminals, as they actually had got some different advices on how to design the physical layout and the colors of their terminals and keyboards from their ergonomic consultants. They belived the black/white screen or the black/green screen to be optimal, but were interested in our results concerning the use of the brown/yellow screen. I informed them about our experiments on anti-glare filters, and promised to send the information on the spray filter, which we had found to be optimum. I also promised to find out whether or not they were able to get a copy of our anti-glare test report.

We then discussed some specific problems and requests concerning the current microbee series. They admitted, that there had been some problems in establishing a sufficient distribution of applications information to the users, especially related to the possibility of user-defined functions in the terminals. We were however welcome to contact Craig Winget, the Supervisor of Special Software, directly. Related to my questions on their power supplies, which occationally seems to show problems on European devices, I was told that the standard 60 Hz power supplies, which were very reliable in the States, were equipped with extra components to handle the 50 Hz requirements. They also had a dedicated 220V/50Hz burn-in line. They had howver once found a batch of power supplies, which had a noise spike occasionally causing trouble, but that problem should have been fixed. We were anyway welcome to contact Jeff Hillstrom, who was the supervisor of power supply products.

I was then shown the factory, which was empty due to the four day working week. Actually, it was much smaller than I had expected, but seemed well organized. I explained about our problem with the RS-232 connectors which could loose the contact with the pcb board, and I was told that I was the first customer reporting that problem. The pcb board was designed for mounting the connectors with screws, but these had been omitted in the production process, as Beehive believed they were unneccessary for normal use. They had however a screw kit, but admitted that this was not generally known.

They did quite amount of testing during the production process in addition to the burn-in, and they believed that this combined with their incoming inspection was the reason for the high reliability of their VDU´s.

In general, the visit did provide a very positive outcome, and the people from Beehive seemed to be very open to all questions.

5´th Conference on Local Computer Networks, Minneapolis, Minnesota.

****

The 5´th conference on LCN´s was held October 6-7 in Minneapolis, and was sponsored by the Computer Center, University of Minnesota and IEEE. The conference committee was amazed by the incredible number of attendees, where more than 160 persons from all kind of companies and countries had registered. This was an indication of the awareness of local computer networks in many areas.

An introduction to local computer networks was given the first morning by H.A. Freeman from Sperry Univac. Although nothing especially new was disclosed, the paper gave a very good overview in which all kinds of local communications was characterized and compared. Especially the taxonomy of local network architectures was discussed in details with examples of networks already in existence. Also the different forms of topologies and LCN´s in connection with gloabal networks was discussed. Also the different activities concerning standards definition, network operating systems design, and future applications was described.

R.L. Gordon from Prime Computers had a good paper on the Perspectives on the Evolution of Commercial Local Networking. One of the reasons for today´s interest in local networking was the fact that more and more companies increased their use of data processing by removing data power from a centralized facility out to the actual users, thereby creating a need for intercommunication between the users terminals and data processing resources. Also, wiring costs were largely labor costs, thus making the intercommunication steadily more expensive and maybe prohibitive compared to the costs of terminals, whereas the price of logic needed for communication and multiplexing was steadily dropping. Thus, the goal was to achieve a very simple and cheap method of intercommunication.

Also the ability to detect failures and dynamically reconfigure a network would become one of the major factors in the future. The trend would also be to put more

intelligence into a node, so that the connection to a network should be even easier than the installation of a telephone. The ability to unplug a node at one place and plug in the node in some other place without notification of a system administrator was highly desirable.

Concerning the use and control of the attached resources to a network, so-called ultra high level protocols in the form of contract protocols would be needed, where an intelligent node actually could put a request for work out on the network, thus enabling resource nodes on the network to make a bid for the job or task, thereby establishing a processing "contract" between a requestor and a resource.

Bob Loveall from Sperry Univac presented an interesting paper on local networks in military applications, where the requirements are quite different from those found in commercial networks.

A. Sauer from Siemens in Germany described a local microprocessor network for the office environment with an optical bus system. The system was an experimental system, using a centralized optical star-coupler or mixer, thus allowing a multiple access broadcast optical bus. The maximum signalling rate was 16 Mbit/second, but this would be increased to 32 Mbit/second at a later point in time. The maximum distance between nodes in the system was expected to be 1-2 kilometers, and a maximum of 50 stations could be connected. Instead of data communication, the network would be able to support 250 voice channels of 64 Kbit/second.

A modified form of the HDLC protocol was used in order to support the specific characteristics of the system.

Brian Shanning from the MITRE Corporation had an interesting paper on Privacy and Authentication for the Automated Office. He defined privacy as the protection against "passive" attacks, meaning that users should prevent release of information and traffic analysis by any unit attached to or outside the network. Authentication was defined as protection against "active" attacks.

He describe the MITRENET and an experiment in developing a secure electronic mail system MEMO (MITRE Encrypted Mail Office) using public key cryptography.

Basil Maglaris from Network Analysis Corporation presented a paper describing an integrated broadband local network architecture.He described the basic differences between long haul networks and local networks, and the need for integrating a variety of communication requirements within a local network. The different architectures of coaxial-based networks was described:

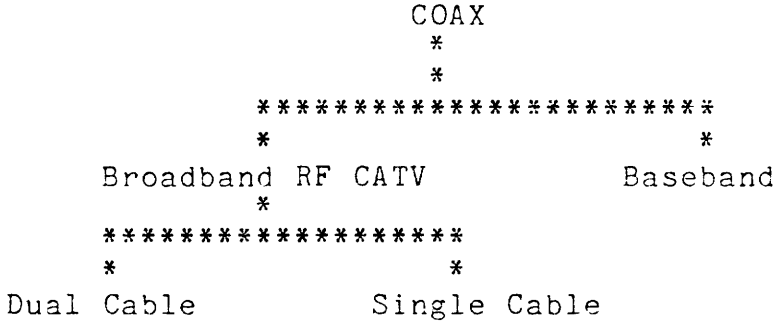

The factors in favour of baseband networks and against RF networks were:

Low loss at lengths less than 1000 feet.
Direct tapping
Passive medium
No needs of modems.

The factors against baseband networks were:

Incompatible with RF components
1 channel limit
Limited distance (3000 feet)
Limited topology (bus)
Low noise immunity
Interference with 60 Hz power lines

The factors in favour of broadband RF CATV networks were basically:

Unlimited distance
Flexible cable layout
Noise immunity to low frequencies

He then described a local network architecture using broadband techniques, where two coaxial cables were used, and where the raw bandwidth was divided into 40-50 standard 6 MHz channels. This would allow different communication methods as bursty, delay sensitive traffic would be served by a contention channel, working much like Ethernet, and where wideband transmission as file transfers, voice or video would be served by several reserved channels. However, a centralized control station would be needed in order to schedule the use of the reserved channels.

In the session on Network Architectures, an extremely poor paper was presented by Paul Kreager from Washington State University Computing Center. The title was No-modem Local Networks, and the interesting research they had made was the observation that 1200 bps asynchronous terminals could be connected to a computer without modems over distances longer than the 50 feet specified in the RS-232C standard, a fact which have been known and used for years by data processing users, including SDC. He also described how to install a diagnostic facility for this kind of network, which was a huge cross switch or junction board, by which the user could direct the cables from the communications controller to the actual cable feeding a specific room, and parallel a monitor cable to a given link.

Through their theoretical research, they had also figured out, that it was possible to put jumpers or straps in the connector attached to a terminal, so the DTR would feed the DSR, thus needing only the send and receive data lines in the cable.

In the session on broadcast systems, W. Bruce Watson from Lawrence Livermore National Laboratory presented an

interesed in prolonged theoretical work, as this would be a hindrance to the standard. All major companies were represented in the committee, and special emphasis was made about the contributions of the semiconductor companies, because it was felt important to achieve an implementable standard, rather than an theoretically optimal standard. The contributions from the companies had been enormous, as all participants had been actively working, both during the meetings and between. A lot of fighting and argueing had taken place at the meetings, and everybody was warned, that they should be able to do a compentent defense of their points of views, if they were going to participate. It was hoped, that the committee would be able to come out with a standards proposal before the publication of the Ethernet standard. By the way, the Ethernet group was also represented in the committee.

The purpose of the committee was to define the media, the DLMAC (Data Link Media Access) and the interface to the network. The concept would be quite similar to Ethernet, but better defined. The committee had chosen the ISO Open System Interconnect architecture, but would stay at the lower two layers. They actually felt, that the division of the lower two layers was not sufficient to cope with this kind of work.

The IEEE network, although basically a multiple access, broadcast bus, was intended to provide transparancy to topology, media, speed and encoding techniques. A modified form of manchester encoding had originally been proposed.

The basic protocol format would contain a sync preamble, destination and source addresses, a control field, an information field and a 32 bit field check sequence, but no length field. Several forms of addresses would exist in contrast to Ethernet, as addresses could be of 16, 32 or 48 bits in lengths. A modified form of HDLC would be used at the bit level.

The interface to the transciever would contain the receive and transmit data lines, RTS and CTS signals, and two error indicators for receive and send errors.

It had not yet been decided, whether or not to include voice communication, and whether or not prioritized messages

should be implemented. Also the use of the CSMA contention scheme had been discussed, because this would require rather long preambles at higher speeds.

The physical prperties would be 1-2 kilometers in length of the network, and the speed would probably be 10 Mbit/second.

Visit at Zilog, inc.
Cupertino, California.
****

The Zilog meeting was prepared in advance by contact through the Zilog European headquarters in Maidenhead, England.

Present at the meeting was:

John MacMurray	Technical Support Manager Components Division
Barry J.R. Barrett	Z8000 Product Marketing Engineer Components Division
Bill Hesley	General Systems Marketing Manager
Jagi Shahani	

First of all, a general presentation of SDC and our microprocessor development activities was given. Zilog had already been generally informed through Zilog in England about our company and our PLZ experience.

We then went on to a general discussion of the Z8000 16-bit microprocessor and the peripheral components in the Z8000 family.

The delivery schedules of the Z8000 was discussed, especially concerning the DMA, which should be in production at Zilog during the fourth quarter of 1981. It is however AMD, who the responsible for the design of this chip. There exist a mask exchange agreement between Zilog and AMD on the hitherto defined Z8000 chips. Further extensions to the family can be developed by each company independently, and are not covered by the current mask agreements.

I was given a general presentation of the Z8000 software activities. The current released software consist of the Z8000 software development package, including PLZ/ASM, the Linker, Imager, up- and download utility. The PLZ/SYS compiler was supposed to be released during the month of

November.

It turned out that the macro preprocessor utility MACP was withdrawn from the Z8000 software package, as there was some inherent errors in it, which in some situation could scratch the entire disc of the development system. I found out, that MACP had been a summer project done by a University student, and that nobody in Zilog could crasp his code nor the very poor documentation.

A disassembler and a floating-point package had recently been announced by a user in the Zilog users' club.

We discussed the debugging facilities for the Z8000, which currently was made up by the ZSCAN 4 MHz real time analyser, an extension module to existing development systems. ZSCAN provides the user with 4 Kword of mappable static RAM, which is protectable as read-only memory. ZSCAN contains break-point facilities, SW trace of up to 240 instruction steps, and is intended as a low cost debugging tool. The cost of Zscan is 4500 dollars.

The next step in debugging tools would be the EMS8000, which would be an 8 MHZ real time analyzer with facilities for symbolic debugging, disassembly, and with its own command language. The price would be around 12000 dollars.

From the outside vendor, MICROTEC, a Z8000 crossassembler (written in FORTRAN) was available, together with a Z8002 simulator. MICROTEC were also capable of delivering a Z8 + Z80 crossassembler and simulator.

The company MULTITECH from Taiwan had chosen the Z80 MCZ software development systems for the provision of crosssoftware to all the commonly used single-chip microcomputers.

Zilog was currently marketing a software package for virtual systems, which included a Z8000 crossassembler (written in assembler) for DEC PDP-11 computers, under the operating systems RT11 and UNIX.

We discussed the efforts related to software support at Zilog, and I expressed the question, why Zilog still had not come out with any released high-level language support for

the Z8000. The answers were manyfold. First of all, the Zilog management did not seem to have paid enough attention to the fact that software support was the key issue to success within the microprocessor business, and had therefore not allocated the needed amount of resources for that. Apparently, they did not quite realize that fact yet. Second of all, there was a huge lack of experienced software people in the Silicon Valley area, especially with knowledge of compiler design, and the job mobility was very high. Some of the very bright people at Zilog had left, either to other companies or starting their own companies. Most of their PLZ experts had left to join the newly formed Exxon company, SUMMIT, which basically was an advanced software development company for Exxon data processing products. Thirdly, the use of high level languages was far less used and demanded in USA that in Europe, and most users of the Z8000 had been satisfied with the initial assembler support.

Zilog would however in the new year announced a new product, incorporating UNIX and the C language, which basically meant that a lot of already available software under UNIX would become available for the Z8000.

There were however also some schedules for releasing PLZ/SYS, FORTRAN and probably COBOL for the Z8000. ADA were looked into, but nothing exactly could be said about this.

Zilog thought, that the biggest problem with PLZ/SYS was the fact that the source-code were not given free at the start. If this had been the case, PLZ would have been the major programming language for microprocessors today in stead of PASCAL, but as the universities could get the PASCAL source-code free and had to pay an immense fee for PLZ, they of course choose PASCAL, which then became fashionable to use in the microcomputer and hobbycomputer world. The PLZ designers at ZILOG had thought however, that PLZ/SYS and its language family concept had been the very best state-of-the-art programming language for microprocessors at that time, and that this fact should be able to market and sell the language easily.

I was presented the new MCZ-2 series of microcomputer systems and the Z-NET concept. As already guessed, the Z-NET was build around the current Z80 SIO chip, which meant that Zilog would be unable to utilize newer technologies with

co-processor. They had however signed a contract with a Canadian company, which would deliver an APL interpreter for their Z80-based MCZ systems, and possibly for their future Z8000 systems. The availability of APL had not yet been announced.

I was presented with some preliminary specs on the Z8065 burst error processor, originally designed by AMD, which is able to do error detection and correction for high speed data transfers at a data rate up to 20 Mbits/second. Also the Z8068 Data Ciphering Processor was described, which is able to encrypt and decrypt data using the NBS encryption algorithm at data rates over 1 Mbytes/second.

Users were however free to specify, and get produced their own specialized co-processor. I suggested the idea of having a co-processor with UV-PROM storage or a piggy-back version like the Z8 piggy-back prototyping microcomputer, which Zilog did not seem to have thought about.

We discussed the implications of the Z8000 control signals on bus architectures, especially the concept of distributed memory managemet units. The questions I asked here were too technical to be answered immediately, so we decided to postpone them to another meeting with some more skilled experts on the actual chip design.

I was also told of the yet-to-be announced Z8003 processor, which would make it possible to implement a true virtual memory concept, as this processor would be able to retry an instruction after a page fault. This processor would be accompied by a new memory management unit.

I then asked about the ZBI bus concept and the new Zilog boards for this bus structure, but it turned out that this series of boards was about to be withdrawn, mainly due to production resource problems. The actual problem of how to announce the withdrawral was still being discussed in the Zilog Management at the moment. Zilog indicated however, that the board series would supposedly be used for internal use in some projects.

During the meeting, we constantly run short of the initial schedule and the planned visit to the wafer production plant were postponed to another meeting.

EF-57 January 1981

Meeting at INTEL Corporation,
2625 Walsh Avenue, Santa Clara.
****

Intel is perhaps the most well-known manufacturer of microelectronics today. They are in many respects an industry leader within the microprocessor field, and this is reflected in their very strong security system in all buildings. I were however able to get inside Intel on two occasions, the first of which was a visit to their litterature department in an attempt to get the Ethernet Specs and specs on the iAPX32. Although these had been announced to be available, it turned out that they were still in press, but I was allowed to take any other needed Intel information, which included some of their just released data books and applications manuals.

The other visit was arranged through Olivetti, and the purpose was to discuss the Intel product strategi for high-end microprocessors, which includes the discussed concept of moving software onto silicon. Attendents at the meeting was:

Serge Rancier Customer Marketing
Nothern & Central Europe

Mike McCullough Product Sales Manager
International.

I gave a presentation of SDC, our TP2 project, our current Z80 designs and our interest in the development of high-end microprocessors and local networks. I was then presented an overview of the Intel Corporation. Intel has currently 17000 employees, the half of which are in the United States. A third of the employees in the States are located in Santa Clara. The annual turnover is 850 million dollars.

Intel is made up by several divisons, of which the most important related to my visit is MSCD (Microcomputer Systems Division) and MCCD (Microcomputer Component Division). The actual organization of these divisions are shown below:

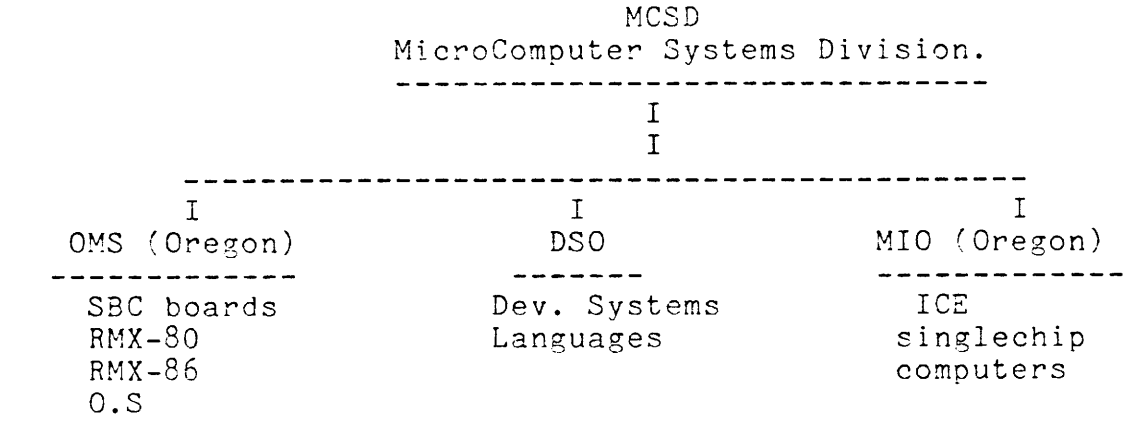

Additional departments exists for Sales, Finance and Legal.

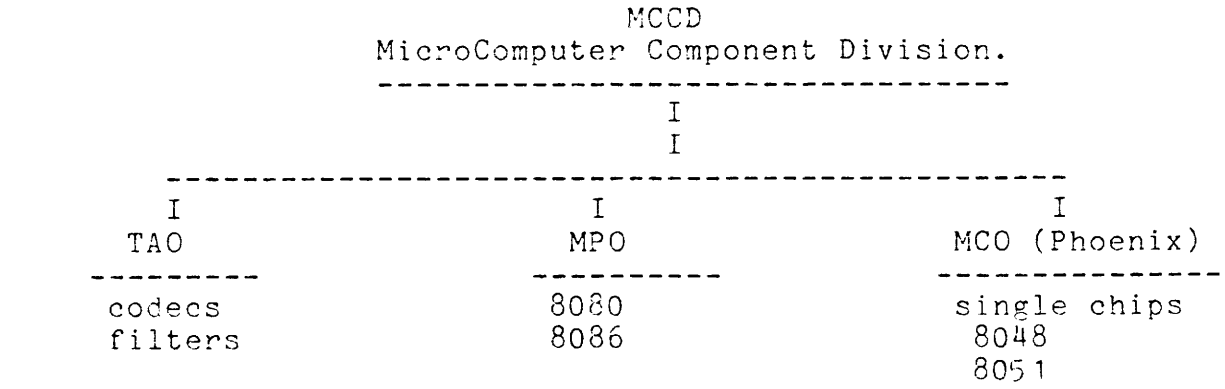

We discussed the evolution of microprocessors related to the software needs, and Intel pointed out that the 8086 originally had been designed for Pascal. They were however well aware, that the marketing of the 8086 had not sufficiently taken this fact into account, but rather sold the 8086 as an upgrade of the 8080. They had plans of increasing their software support in the future, as this were still a dark area in the Intel support of microprocessors. They were beginning to pay special attention to software related problems when designing the architectures of microprocessors, and had come to the belief, that the reliability and ease of writing software were more important than pure performance. As one specific

example, they choose to use a stack frame for parameter passing during procedure calls rather than passing values in registers.

We discussed the architecture of the iAPX 286 processor, which is designed for a more object oriented processing. Also the different forms of memory management were mentioned, where Intel believed that the use of a very few set of registers were more attractive than the Zilog approach, especially as this would allow an easy integration of memory management on the CPU chip.

Related to the iAPX432 32-bit processor, Intel were designing this in Oregon, and the Santa Clara Divisions had not much detailed information on this yet. We talked about the problems of guaranteeing the quality of the software put onto the chip, and Intel supposedly thought, that their experience, especially with the RMX operating systems, would justify this approach.

Concerning Ethernet, I was told that Intel were in progress of designing some interface chips, but that nothing actually had been put onto silicon yet. They had the feeling, that the Ethernet was sufficiently simple and powerful to become widespread used in the local networking envirronments, so Intel together with the other two partners, Xerox and Digital, had decided to push the concept in an attempt to make it a de facto standard, rather than waiting for a standard committee to agree upon something, which might only be a minor enhancement of Ethernet.

We discussion the possibilities of establishing some non-disclosure agreements, especially on the iAPX432 ADA processor, and I was informed that this would be no problem for a company like SDC, but that the arrangements had to take place through Scott Dixon, Intel, Bruxelles. They told me, that by the time I would return to Denmark, I should be able to get the Ethernet Specs through Tom Stenberg, Intel, Denmark.

Meeting at Fairchild,
464 Ellis Street, Mountain View.
****

Olivetti had suggested, that I should visit Fairchild as a part of my study tour. I had therefore arranged a meeting at Fairchild Camera and Instrument Corporation, 464 Ellis Street, Mountain View. The agenda I had suggested included the Fairchild developments within peripheral components, bipolar support circuits for microprocessors and the integration of on-chip test facilities on VLSI components. At the meeting attended:

Hank Miller	senior appl. engineer bipolar LSI
John Le Gall	bipolar LSI QA manager
Valerie Millar	customer relations coordinator.

It turned out, however, that Olivetti had not given them a qualified information on our company and the reasons for my visit, so the people at the meeting were mainly working in the fields of semiconductor test equipment and very high speed bipolar bit slice processors. They were anyway able to get hold of some other people, able to discuss my questions. We also arranged, that I should visit some of the other Fairchild divisions in South San Jose and Palo Alto, where the actual work in microprocessors were taking place.

Anyway, we got a good discussion on the progress within the bipolar area. Fairchild had supposedly no plans of offering the kind of TTL support circuits like Texas Instruments. oriented towards the requirements of new microprocessor bus based systems. They had the opinion, that the integration of self-test capabilities and redundancy at the chip level would not be justified by any means. Actually they did not seem to take the users' needs for rapid failure detection and recovery into consideration in their chip design. Their bipolar activities were oriented towards the use in bit-slice processors, where they argued that in such

systems, the users would themselves add the needed hardware to ensure failure detection.

Meeting at Fairchild,
101 Bernal Road, San Jose.
****

This meeting was arranged in order to get some information on Fairchild's developments within the area of communications chips. Advanced communication IC's may perhaps be one of the most significant activities of Fairchild related to the microprocessing world. At the meeting attended:

Roy S. Kole	Manager Engineering
William R. Nienhaus	Microprocessor Peripheral Design, MOS special products
Exkhard Dornbrach	International Marketing Manager, MOS Microcomputer Group

I explained SDC's objectives of having interest in the progress of semiconductor communications technology, especially the impact of communications technology on future system architectures. We then went on with a discussion on specific Fairchild parts and the trends of communications technology in general.

Fairchild is current designing peripheral communications chips aimed at the F6800 microcomputer family, which is a second source of Motorola's 6800-series. Most efforts go into the design of the so-called advanced data link controllers, capable of handling the ADDCP, HDLC and SDLC protocol standards. One of the chips in this series is the F68A54 ADLC, a 5-Volt only chip capable of maintaining HDLC/SDLC communications at a 1 Mbit/second data rate. Another more recent development is the F3846/6856 synchronous protocol communications controller. This chip can handle the previous mentioned bit-oriented protocols SDCL, ADCCP and HDLC together with a full support of byte control protocols like IBM's BISYNC, DDCMP and others. It supports data rates up to 1 Mbit/second and is programmable to handle the byte oriented protocols in either ASCII or

EBCDIC mode. As common for the new generation of communications chips, it contains internal self test capabilities. It can be programmed to use either an 8-bit or a 16-bit data bus and is thus oriented to the new 16-bit microproessors. An enchanced chip, the F6856A, which would include an asynchronous controller was scheduled for production in 3´rd quarter 81. This chip, as the F6856, would contain only one channel, and the Fairchild people thought that the Zilog SIO with roughly the same facilities would be a strong competitor. The Fairchild chip would however permit a more easy interface to the 6800 line and contain the full support of the IBM BSC protocol in contrast to the SIO.

We then continued with a more detailed discussion of the actual architectural design of the chip, by which I pointed out, that a chip like the Zilog SIO was more suited to the new generation of microprocessors, as the chip actually contained the advanced capability of interrupt vector generation, which would be an enormous advantage, as the 16-bit processors like 8086 and 68000 had recognized the importance of a true vectored interrupt capability, and as the capability of true vectoring would result in a great simplification of the task of writing interrupt software. Also the need for a separate interrupt controller would be obsolete, making the use of the communication chip a more cost/effective solution. It seemed however, that the Fairchild people had not much software experience, and was not able to fully grasp the interaction between chip architectures and software architectures.

The future work of the communications group at Fairchild was concentrated on developing a chip supporting the X.25 level 3 protocol, like the announced Western Digital chip. It was still under consideration, whether this chip should also support a complete SNA link or if a SNA controller should be implemented in a separate chip. This led to a discussion of the trends in communication chips concerning the integration of several protocol types in one chip, which possibly could impose some restrictions on the amount of support of each protocol, or the provision of very higly supported protocol, each in a single chip, but possibly with the same pin-out.

Concerning Local Networks, it seemed that the people present at the meeting had no knowledge of this area, and they could

therefore not give any information on Fairchild's future plans for this concept. This can be caused by the fact, that the develoment of local networks includes the need for some software expertise too, and that the high speed of local networks actually makes the chip design a bipolar problem, located in some of the other divisions of Fairchild.

The need for data encryption in our networks were then discussed. I was presented the initial specs for the F6857 data encryption circuit, a chip which would first be launched in 1st quater 1982. This chip will meet the requirements of FIBS Pub. 46 - Federal Encryption Standard and would contain an on-chip DMA controller, making the actual task of en/decrypting a message very easy. The data rate, however, would be only 140 Kbytes/second in electronic code book or cipher block chaining modes, and 24 Kbytes/second in 8-bit stream cipher mode.

Finally, we had a general discussion of the microcomputer market and Fairchild's activities. Fairchild was currently evaluating the announced 16-bit processors in order to establish a second source agreement on one product. I explained our reasons for selecting the Z8000, and pointed out the basic problem of software support, not only for this specific processor, but for the entire microprocessor market in general. I noted, that were they to choose a processor currently taking care of the problem of providing capabilities of reliable software, they should look for the NS16032 from National Semiconductor, as this processor had a seemingly well designed instruction set, although it had not yet seen any silicon. Probably, the software problem would anyhow be solved for the Z8000. The people at the meeting seemed to have a very good knowledge of Fairchild's plans concerning second-sourcing, but they would not give any kind of hint. We also discussed Fairchild's range of single chip microcomputers, the F3870 and F8 series, and the software support of this.

The basic impression was, that the Fairchild people was competent on the technical details of data communications, but that the architectural overview had its weak links.

Meeting at Fairchild,
4001 Miranda Avenue, Palo Alto.
****

This meeting had been set up to give me some information on the advanced microcomputer development at Fairchild, especially within the 16-bit area. The schedule was very thight, as I came directly after noon from the ACM conference in San Francisco and had another meeting at Zilog in Cupertino later that afternoon.

I basically talked to some of the people involved in the design of the 9445 Microflame II 16-bit microprocessor, including:

Linda McCleary	9445 Marketing
Chuck Ericsson	9445 design engineer
trong-Tich Dao	Manager, Advanced uP. Systems

The 9445 Microflame II microprocessor has been developed as an outgrowth of 9440 microprocessor, with an upgraded instruction set and more advanced architecture. It is implemented using Fairchild's Isoplanar integrated injection logic technology, I3L. Basically, it is a single chip equivalent to the Nova minicomputer.

As being a minicomputer equivalent, it supports operator console functions including internal self-testing. The instruction set and architecture are also impacted by the origin, and leaves much behind compared to e.g. the 68000 and the Z8000. The addressing capabilities restrict themselves to 64 Kwords or 128 Kbytes.

I discussed the reasons for Fairchild entering the design of this type of microprocessor and the possible future evolutions of the family. It turned out, that the basic reason was the amount of software already written for the Nova computers, and that Fairchild thought that a single chip version of the minicomputer could have some applications. There would come no major architectural

inventions, basically because of the restriction in memory addressing and the software support. Fairchild had designed some support circuits including a bus arbiter and and dynamic memory controller for the 9445 family, but even though I was presented the diagrams for these, they had still not been submitted for actual mask and chip design.

There was no plans for providing any other peripheral support circuits. Fairchild supposed peripheral controllers to come from second sources, belonging to another microprocessor family, like the 6800 or 8080.

We had a specific discussion of the possible future architectures of microprocessors and the way e.g. software architecture would affect the hardware architectures. Fairchild admitted, that they would first of all go into some kind of second-source agreement on announced microprocessors, but that an entire new generation of advanced microprocessors from Fairchild would possibly be an ADA machine, although they had no participation in the ADA work, which has been carried out. This is an remarkable fact, as there seems to be an increasingly need for semiconductor manufactures to take part in standards definitions in order to get implementable standards.

Meeting at Ungermann/Bass Inc.
2650 Mission College Boulevard, Santa Clara.
****

This meeting had been arranged in advance from Denmark. Ungermann/Bass Inc. is a very young company, established a year ago by Ralph Ungermann and Charlie Bass. They are best known as the inventors of the PLZ/SYS language at Zilog. Getting tired of the internal organization at Zilog and the problems of getting their ideas through, they left Zilog to form their new company, which in the short time since the start has become one of the leading companies within the area of local computer networks. Their first product is the local network concept NET/ONE.

At the meeting, I met:

Charlie Bass	President
James F. Jordan	Vice President of Marketing
John M. Davidson	Senior Network Architect

I started with a general presentation of SDC's current online system, the architecture of our new TP2 system, and our interest in both local networking and public data networks. This was a subject of great interest to the Ungermann/Bass company, and we had a long discussion on specific details of our network and operations.

I was then given an introduction to the company, the product range and the trends in local networking. According to Charlie Bass, a specific trend was clear: in the 1960'ies communications technology was used in two areas: Long-distance, where intercommunication was established between remote mainframes on one hand and the connection of terminals to remote computers on the other hand; Locally, where intercommunication was established between local mainframes on one hand, and the connection of terminals through I/O-controllers to local mainframes on the other hand. In the 70'ties, public VAN's were being used to establish a common intercommunication interface between system components, regardless of these being computers or

terminals in the case of long-distance networks. Locally, the concept of local networks were now being established to provide a common intercommunication interface between system components with a fully distributed control, independently of the actual nature of the components, i.e. terminals or computers. In the 1980'ies, the trends of VAN technology and local network technology would probably merge into a state of so-called total processing and intercommunication.

The specific features of local networks could not be stressed heavily enough. They had the distinct feature of being geographically restricted, thus enabling a very high speed communication over a relatively simple medium, thus offering a significant improvement in reliability, and finally offering a hitherto unseen possibility of privateness, meaning that security control could be managed within the network, no official regulations would probably not affect the use of the network and the administration of the network would be extremenly simple.

The basic feeling was, that the shared, multiple access, broadcast bus like Ethernet or the coming IEEE standard would be the concept best suited for the range of applications found in everyday data processing. Several other architectures for local networking had been proposed, but these did not have the elegance and simplicity of the broadcast bus. One interesting fact, which had been noticed in the States, was that the topology of a given local network selected in a company actually would reflect the organization, which they were serving. That meant, that highly centralized organization would probably not be the first to use local area networks, as the kind of distributive control in such a network would be an unconscious treath against their internal organization.

I was then presented the concept of their NET/ONE local network. NET/ONE is actually not a "pure" local network in the sense, that individual terminals can connect directly to the cable. Instead, they have chosen to develop an interface box to the network, allowing any type of terminal or computer to connect to the box through any type of protocol and thereby having access to the network. This means, that standard RS-232 compatible terminals or BSC terminals can be supported, and that all protocol conversion and data conversion (.e.g. from EBCDIC to ASCII) takes place in the

box. The box itself is a Z80 based microcomputer, managing the network link and the link to the external devices. The entire systems programs are programmed in PLZ/SYS.

One of the basic areas of use is within the development systems integration, where a user can employ the NET/ONE to integrate the access to different development systems or host computers from a single terminal. Thus a programmer is able to establish a link from his terminal to e.g. a Zilog development system, perform some work on it, then establish a link from the Zilog system to e.g. a minicomputer and initiate a file transfer, then disconnect the links, and transfer his terminal into an APL-terminal by establishing a link to an APL-processor etc.

Another area is the interconnection between systems in a computer room, where most equipment has some sort of communications interface, but generally not a dedicated local networking capability.

I saw the development laboratories in the company, and was given a demonstration of the actual use of the network within the company ifself. We also discussed the future trends in local area networking, and they had the opinion, that there was no need to sit and wait for the standards to come, because the potential in the industry was already large enough to justify this type of network, which anyay could be modified to give access to e.g. the Ethernet or the IEEE net.

We had a specific discussion on the special coaxial connectors used to permit an easy attachment to the coax cable, and I was provided with a sample.

Finally, we discussed PLZ/SYS and the developments of the language. Charlie Bass noted, that most of the PLZ people at Zilog had left, either to his own company or the newly founded Exxon-company, SUMMIT. He had previously had some contact with Olivetti on the language, especially concerning the establishment of a PLZ users' group, but nothing had happen. I was recommended to contact Mike Fay at Summit, where he felt the most developments in PLZ were currently being made.

Visit at AMDAHL Corporation,
1250 East Argues Avenue, Sunnyvale
****

This meeting had been aranged beforehand through Amdahl A/S, Denmark. The topics of interest was the design and use of bipolar LSI and VLSI technology in Amdahl computers, the possible architectures of large CPU´s and the use of CAD/CAM technigues within Amdahl.

During the meeting, I met:

William F. O´Connell, Jr.	Senior Vice President International Operations.
Bill Fry	Marketing Support Representative
John Foggiato	Process Technology Manager
Mike Clements	Vice President, Engineering
Herp Hellermann	Principal Planner
Steve Kardys	CAD Manager

First of all, I got a general introduction to Amdahl and their current product line, together with a lot of sales information.

I then gave a presentation of SDC, my project and the background for our interest in VLSI technology. After this, I was presented with an overview of Amdahl´s currently used chip technology, their process technology and their concept of multi-chip carrier boards. We made a visit to their LSI Laboratory Display, a show room in which they had a selection of wafers during the various production steps, together with a display of the entire chip production process. There was actually no specific news in this, but the show room may have more interest for people not accostumed with LSI technology.

We then made a facilities tour, where I was shown the chip fabs, the MCC board assembly lines, the mainframe assembly lines, Gene Amdahl´s very first computer with a bullet hole from a shotgun and their software development division.

Afterwards, we had a technical discussion of the evolution of large CPU´s and the technology behind. I was given Amdahl´s opinion of the state of the technology by 1985. They felt, that by 1985 GaAs technology would have a major impact on the data processing industry. The needed chips for the CPU in a V-8 sized computer by 1985 would amount to 20-40 chips. The equivalent number of gates in NMOS technology would be implementable on 6 to 13 chips. It was expressed, that the computer design of 1985 - 1988 would give a 30 MIPS instruction rate and a data rate of 2000 Mb/s. They could also foresee a development towards on-chip cache for bubbles in 1985.

The Amdahl views on communication was very traditional for mainframe applications. They have accuired the communications controller company, TRAN, which enables them to provide their systems with IBM-like communications controllers. They had no ideas concerning the progress of local networking nor local mainframe networks, and could not envisage the potential need for a local mainframe network. They were however aware of one standards effort on high speed link networks, namely the ANSI X3T9.5 LDDA (Local Distributed Data Architecture) standard.

They were investigating the use of fiber optics as a replacement for cabling within the CPU in future mainframes. Related to the need for larger addressing spaces in modern applications, supported by the advent of cheaper memories and possibly bubble technology, they might enhance the addressing capability to 31 bits, but not more, mainly due to software problems. Actually, they had a modified Amdahl V7 running with a 31 bit addressing range in the laboratory. The keyword of future large CPU´s was expected to be performance, and the ability to change or reconfigure in correspondence to changes in functional requirements and changes in data flows.

Finally, I made a tour to their CAD/CAM division, where I was presented for their CAD applications. It was interesting to see, that the actual use of CAD systems were very low.

All chips were laid out by hand and pencil, and first when the actual drawings had been completed, they were digitized. It was possible to make minor changes interactively at the CAD system, but in general, the entire chip had to be redrawn by hand in the case of a major change. However, the number and the complexity of the chips used in the Amdahl systems are rather low, justifying this approach. Memory chips and certain bipolar control circuits are brought from outside vendors.

Pcb design are done interactively through an Applicon equipment, but colors are not used to discriminate between the numereous layers of the complex MCC boards. The actual boards are manufactured by outside subcontactors, and the pcb masks for the multi-layer boards are mailed on photographic films. They did not seem to have any problems with the stability and aligment of the films, although they admitted that they had to perform a very cumbersome procedure in order to control the humidity and temperature stresses of the films during mail.

In general, the visit gave a good introduction to the production process of a large mainframe manufacturer. All the people involved was very competent and relatively open, although no information not generally known was given.

EF-57 January 1981

Visit at Western Digital Corporation,
3128 Red Hill Avenue, Newport Beach.
****

Western Digital is best known for its line of data communications circuits, floppy disk controllers, and the PASCAL-microengine, a microprocessor capable of executing p-code directly. The meeting had been arranged through C-88 in Denmark. The day selected, however, turned out to be the day before Thanksgiving, which caused the entire company to be occupied by the thanksgiving preparations (all managers were busy by provided the employees with turkeys etc). Also some of the people, whith whom I should have met, were not present.

At the meeting attended:

Ralph Bubak	Manager, Headquaters Sales
Charles A. von Urff	Vice President and General Manager, Telecommunications Division

I gave a presentation of SDC's current online system, our new TP2 project and the communications aspects involved herein. Also I explained our interest in up-to-date know-how on communications technology.

We then discussed Western Digitals range of products within the communications area. Western Digital has established an internal commitment to provide support for all types of data communications. They have today a fairly good range of advanced communications controllers, especially for SDLC and HDLC.

Of special interrest was their announced X.25 packet controller chip WD2501, which was supposed to come out shortly. The actual mask design had been made, and the chip was now in prototype testing. The chip takes care of the lower two layers of the X.25 protocol, and is capable of automatic retransmission and sequencing.

Western Digital expressed, that the availability of

encryption and decryption in communications networks would be heavily stressed in the future, and they had already some data encryption devices in design.

We discussed the problems of having support circuits and controllers for local networks. Western Digital was of course represented in the IEEE committee, and they were positily sure, that the type of broadcast, multiple access bus, would be common. They were however not quite sure, that the final form would be 100% Ethernet. In fact, they would probably come out with their own standards proposal on local networking, which would be token based.

They felt that the time right now was not ready for voice communication through the local networks, but were anyway confident in that voice would be integrated in all local networks by the late 1980´ies.

Also they recognised the need for a very cheap local network interface. The current price for some local network interfaces was around 12000$, 500$ for the software and 340$ for the cable, a price which they considered prohibitive.

They were not quire sure, whether they would produce an Ethernet-chip capable of being reprogrammable to serve the IEEE standard or their own standard proposal, or build separate devices for each standard.

On the question of upgradability of local networks in response to the progress of technology, they answered that local network architectures should be speed independent and adaptable to use new technology, e.g. fiber optics. They felt quite sure, that this could be achieved, although it was not an integral part of the current standards proposals.

I then asked about their views on software and the impact of software on hardware architectures, especially with references to their PASCAL microprocessor which was an example of a piece of hardware designed from the requirements of software. They said, that PASCAL was definitely out, and that the design of the Pascal microengine had originally been a design exersise in order to see whether or not Western Digital was able to design microprocessors. They were however committed to continue the support of this product. They were anyway working on an ADA

machine, which still not had come to the stage of silicon design. A man called Dave Fisher was responsible for the ADA design.

We then discussed the activities within the disk controller area. Western Digital was currently designing a Winchester disk controller, and they would also provide some controller boards for winchester drives. We then entered a long discussing of why Western Digital and other board manufacturers were not able to design small circuit boards comparable in size to the small sizes of 8" and 5.25" drives. I felt that it was illogical to have a huge circuit board two to three times the size of the disk drive, but WD did not seem to fully realize the need for smaller controller boards.

Western Digital was anyway interested in our views on the applications of the technology, and invited us to join a non-disclosure information exchange agreement, although this was not common between manufactures and end-users.

Olivetti Magnetic Peripherals.
****

One of the divisions at Olivetti, Cupertino, is their magnetic peripherals group.

The major area for this group is the acceptance and evaluation of magnetic peripherals and magnetic components like tapes, heads and platters from outside vendors.

They are however also doing quite amount of development projects, both by their own and in cooperation with Olivetti in Italy. The current development activities are concentrated on rigid discs. IBM was supposedly currently doing major developments within the floppy disc area, but Olivetti had chose to wait and see what exactly happened.

One of the greatest matters of concern was the read/write heads in the drives, and Olivetti holds some patents on head technology.

Olivetti is currently designing their own 8 inch winchester drives with a capacity of 25 Mbyte. The current state of recording technology is providing 500 tracks per inch and 6400 bits per inch for production devices, but would probably improve to 900 tracks per inch and 10000 bits per inch within the next 3 years.

The heads for the 8 inch winchester drives are produced by Olivetti themselves, but platters are brought from outside vendors like AMC (Applied Magnetics Corporation), National Micronetics and Infomac.

For the 8 inch drives, the 3M cartrigde streamer will possibly be used backup. The streamer will be delivered fra DEI.

Olivetti was a little bit afraid of discussing the concept of vertical recording with me. However, it came out, that they felt vertical recording technology for tapes and discs would be 5 to 10 years from now. The problems related to vertical recording technology was not the heads, but the media, as tapes and platters would have to be manufactured differently than today.

Olivetti was also a bit afraid of disclosing their activities within the design of a 5.25 inch mini-winnie with 30 Mbytes of storage. This unit would probably be available within 2 years.

EF-57 January 1981

Visit at Dataquest. Inc.
19055 Pruneridge Avenue, Cupertino.
****

Olivetti was very impressed over a research and analysis company, called Dataquest. By arrangement through Olivetti, a meeting was arranged with Dataquest, who should be very informed on the state-of-the-art within the technology, and who should be able to provide some names on persons of interest at various semiconductor companies.

I spoke with

Daniel L. Klesken	Vice President, Director, Semiconductor Industry Service

However, the outcome of the meeting was not very good, as the person with whom I talked did not seem to have any general knowledge of the area at all. It seemed to me, that Dataquest´s job was to gather statistics and sales figures, and then present these figures to costumers interested in following the market. Therefore I was only able to get some names on sales people, but not on any technical experts. Even though Dataquest was supposed to make evaluations of minicomputer systems like the Datapro services, the reason why I asked the question of whom to contact concerning microcomputer software, the answer was that Dataquest knew nothing about software at all.

The general impression was, that in the United States, companies like this can actually sell a product news service without pocessing any competent technological know-how.

Visit at Memorex Corporation,
2720 Orchard Park Way, San Jose.
****

This visit was offered from the Memorex exhibition stand at the Mini/Micro´80 conference in San Francisco. At Orchard Park Way Memorex have established new production facilities and research laboratories for their OEM 8 inch winchester drives.

The visit was a tour round the production facilities. First of all, the current range of 8 inch drives was presented, together with examples of outside vendor controller boards from Data Technology Corporation and file systems from Cyperdata. Then the actual production area was shown, including assembly lines, clean rooms for platter test and assembly, circuit board assembly, test and quality assurance. The production had recently been started, and the current throughput was only 25-30 drives per week although this would increase drastically during the following months. Memorex was producing all parts of the drives themselves. Approximately 200000 square feet was dedicated to the 8 inch disc drive development and production. The unit in actual production was the 101 with 11.7 Mbytes of storage. An improved version, the 102 with 23.4 Mbytes of storage would be in production by next year.

After that, we saw the development laboratories in a near-by building. It was interesting to note that Memorex actually was using Zilog development systems for the software development for the disc drive controllers. Also shown here was a prototype of the 201 fixed/removable 8 inch disc drive, a drive using winchester technology to achieve 25 Mbyte of storage divided on 12.5 Mbyte on fixed platters and 12.5 Mbyte on a removable cartridge. By providing a removable, dust proof cartridge, Memorex would hereby solve the inherent problem of back-up on high density winchester drives. The drive was of the same size as a floppy disc unit, and was especially oriented towards desk top office and small business systems operating in an office envirronment. The unit included internal self-diagnstics in order to reduce maintenance costs.

The Memorex people was very open, and even allowed

photography in the production facilities and in the research laboratory.

EF-57 January 1981

Visit at Memorex Corporation,
San Thomas at Central Expressway, Santa Clara.
****

This visit had been arranged through Memorex in Denmark. The purpose of the meeting was to be informed of Memorex' views on the trends within disc technology related to small systems, the plans of Memorex within the area of VDU terminals, and the impact on VLSI on communications technology. Unfortunately, Memorex in Denmark did not quite understand the technical words in my agenda proposal, so Memorex in Santa Clara had some trouble in providing the revevant people. They had howver been informed, that SDC was a large IBM/AMDAHL user, so they had prepared a presentation on their large disc systems. They were anyway able to get hold of qualified persons during the day of the meeting.

At the meeting attended:

Fred Jakolat	Manager, Product Planning Storage Systems Group
G.R. Kraft	Manager, Marketing Administration Peripheral Equipment Group
Rolf A. Bertler	Manager, U.S. Relations Memorex Europe Limited
Claire E. Good	Product Marketing Specialist Terminal Products
Thomas J. Quinn	Product Manager 8 inch Rigid Disc Products Peripheral Equipment Group
Shaleen Cole	Guest Relations

I was presented an overview of the current range of Memorex Storage products related to mainframe applications, and with some of the epectancies of the future developments.

The current trend is the move towards thin film media. Memorex had invented a so-called sputtering process for

coating the platters. This was needed in order to achieve the higher desities on new drives.

Related to architectural advances in disc technologies, Memorex was aware of the fact that multiple functions hitherto only found inside the mainframes would be moved to the disc system, thus providing the users with intelligent disc systems or dedicated back-end data base processors. They would leave the variable length block, as fixed block format would be the future approach due to transportability reasons.

The current demand of disc storage was estimated to be 50 million megabytes of online storage, and Memorex expected this demand to increase to 100 million megabytes of online storage by 1984, thus preparing the way to increased market opportunities for both Memorex and other vendors.

Memorex stated several times, that the trend was towards the fixed block format, back end processing, and that the high-end part of the market possibly would be impacted by the small winchester drives.

We then got into a discussion on specific details on recording technology. The current bit cell size at 6360 bits per inch recording density was 156 microinches, which corresponds to 4 particles. At a recording density of 16000 bits per inch, which was anticipated to occur in the future, only two particles will constitute a bit cell with increased risk of noise, thus requirering more sophisticated controllers.

The current technological limit was a storage density of 50 Mbit per square inch, but this could be increased to 100 Mbit per square inch using vertical encoding technology.

Vertical encoding is, according to Memorex, 2 to 3 years away. They believed that the manufacturing of media would be not problem. In fact, they had already been able to perform a test production of medias for vertical recording, and had vertical encoding devices running in their laboratories.

I was then presented the Memorex VDU line. Memorex was very interested in SDC's view on ergonomy, although they claimed to have taken every fact into consideration in their own

VDU´s. Related to VDU architectures, they stated that they were still an IBM follower in respect to functionallity. They were working on color VDU´s as well, and believed to be able to make these of a much better quality that IBM´s.

Concerning the recent trend of IBM to provide ASCII RS-232 compatible terminals, they had no specific views.

I was also shown the production facilities of the hard disc assembly and communications controller division.

Visits at Olivetti Advanced Technology Center,
10430 De Anza Blvd., Cupertino.
****

Although there had been some problems in getting acceptance from Olivetti concerning a prolonged visit at their Cupertino-based technology center, the arrangement was partly made before I left Denmark. At the agreed date, I arrived in Cupertino on the address provided by Olivetti in Ivrea. It turned out, however, that Olivetti had moved long ago, and the other people in the buildning were not quite sure to where, although some of them thought that it was to a bigger two-stories building in the neighborhood. I found the nearest two-story building which of course contained no signs of Olivetti, but luckily enough it was the location of the Cupertino Community Trade Chamber, who was very helpful in providing me with the new address and instructions to find it.

I was supposed to meet Enzo Torresi, the president of the Cupertino division, but I had been unable to get hold of him by telephone before the visit. I found out, that he was away on business trip, but they were expecting me anyway at Olivetti. I was met by Antonio Bartocci, the director of the systems design group, who gave me an introduction to the technology center and its activities.

The Olivetti Advanced Technology Center was between 7 and 8 years old. It had been established to follow some costum design for Olivetti at the semiconductor manufacturers and to handle the shipping and receiving of U.S. components. Finally it was responsible for a certain part of the incoming inspection, especially for the more complex devices.

Approximately one year ago, it had been reorganized to embody dedicated development activities and to maintain a much more closer contact with the Silicon Valley industries than hitherto. The number of employees had increased from 20-25 people to 100 people. As an example, the hardware group now contained 46 design engineers, and the O.S. group contained 12-15 engineers.

The following schematic shows the part of the organization

of interest:

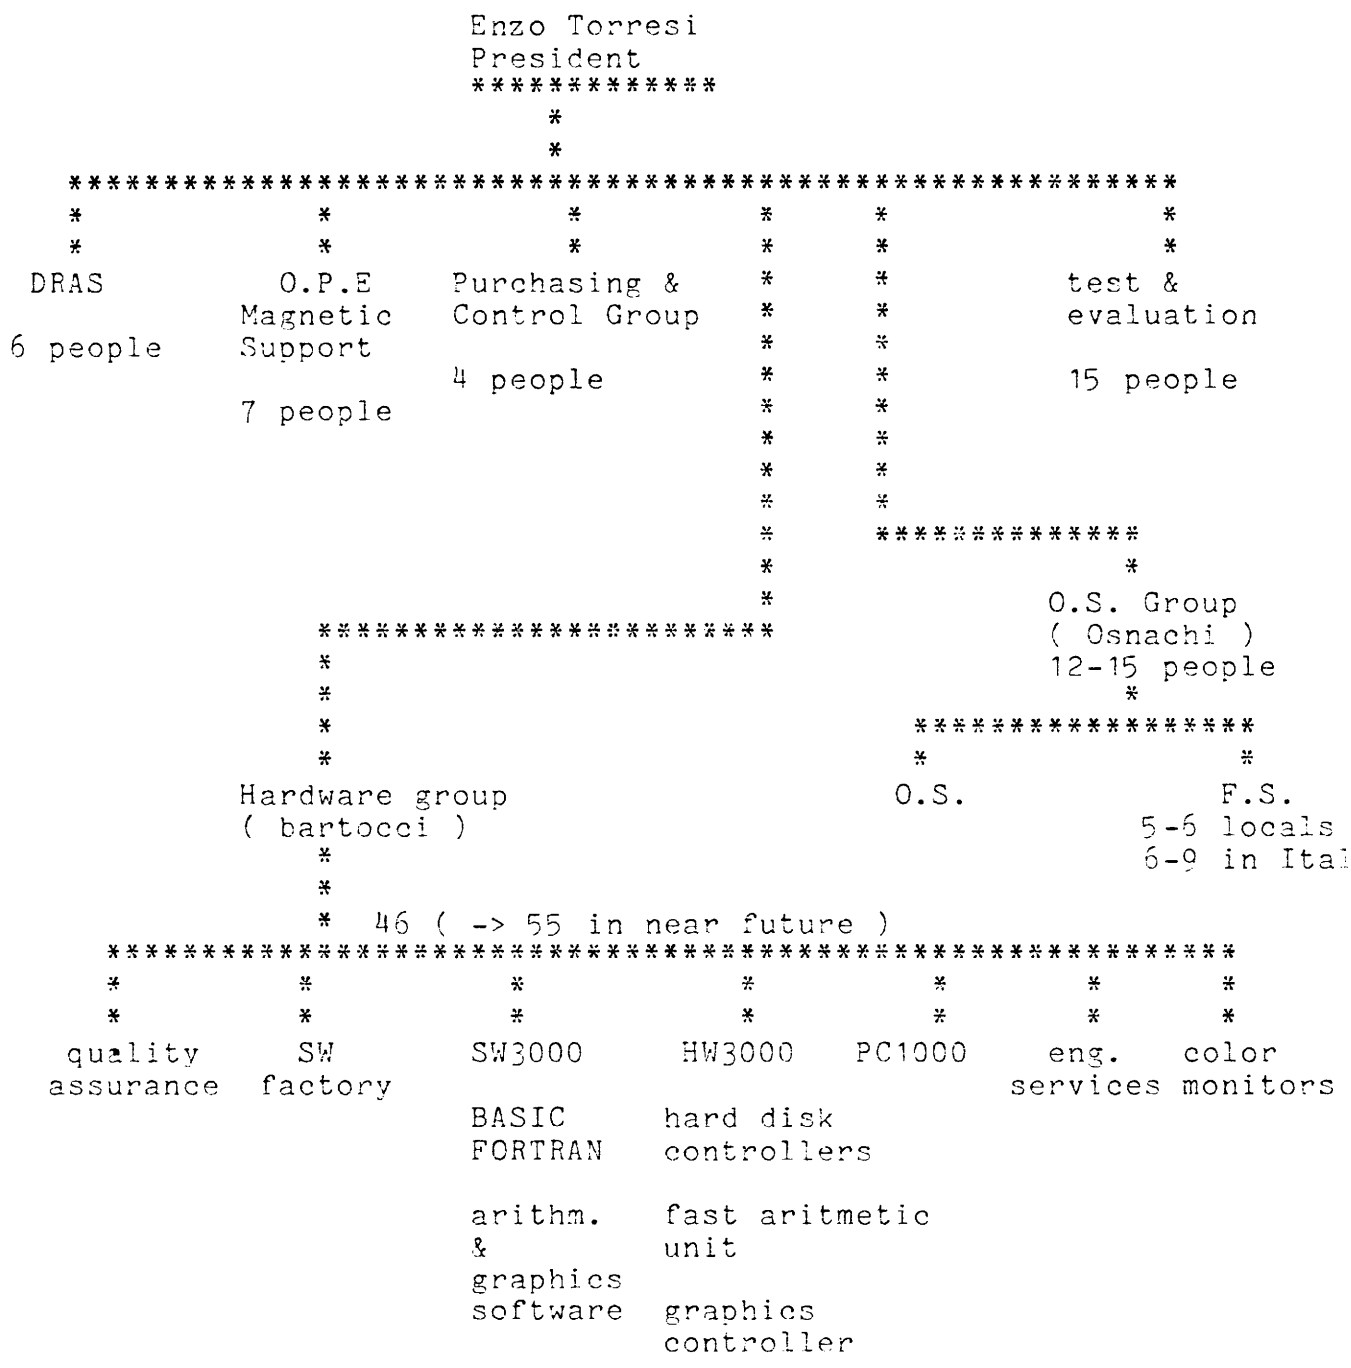

The PC1000 was an unannounced personal computer, intented to

fit in the gap between pure hobby computers and current small system office computers. It would have a price in the range from $3000 to $7000. It would be based on the Z8000.

Bartocci suggested that I should try to get arranged visits at National Semiconductor, Rockwell, Motorola, Fairchild and possibly Intel. He belived, that Olivetti could help to arrange the visits, although they were currently not very popular at Motorola due the rejection of the Mc68000.

I was allowed to come and go at the center, but I should make arrangement through Enzo Torresi or himself if I wanted to speak in more details with the people at the center due to the confidentiallity reasons. I was also given some hints of whom to talk with concerning the different matters of interest:

SW factory:	David Helms
PC1000:	Richard Parker
Engineering services:	Les Dale
Color CRTs:	Mike Swatek
HW3000:	Tom West
Test and evaluation:	Roberto Grasso
DRAS:	Lucas Cafiero
O.P.E.:	Alighier Galvagni
O.S.:	Ashok Saxena Pierluigi Fiorani
Bartocci´s secretary:	Stella Kelly
Torresi´s secretary:	Gabriella Timpanaro-Perrotta
Purchasing:	Dave Birlem

The Olivetti Technology Center had heard a lot about the

"Danish Bank" and the big order, and was very interested in getting a more detailed knowledge of SDC. I gave a general presentation of SDC to Antonio Bartocci, and explained the background of my project and our interest in obtaining know-how within the microelectronics field.

We agreed to postpone further meetings to the following monday, when Torresi would be back. On my question on how they actually provided a living place for short term employees from Italy, I got a lot of help from Enzo Torresi's secretary to find an appartment to live in for the next month, which was both nicer and cheaper than my current motel room.

On the following meeting with Enzo Torresi, I gave a presentation of SDC and the TP2 project once more. He then explained about the technology center in more details and showed some examples of integrated circuits which had previously been designed for Olivetti at some of the semiconductor companies. These circuits had been used in calculators and small terminals, but since the advent of microprocessors and large scale integration, Olivetti had stopped this activity.

We then had a long discussion on the development of LSI and VLSI circuits and the semiconductor companies in the Valley. He believed that the major factors characterizing the different companies was their ability to master either design, especially through the use of CAD, or process technology. Some of the companies were pretty good at design, but had not sufficient control over the process technology, like Zilog, whereas others, like Intel, had a very good process technology but actually did their design by brute force and many people. he felt, that the competition from Japan would be felt in the years to come, as some of the Japanese companies, especially Nippon, were mastering both the CAD techniques and the process technology. What was currently missing was their ability to do advanced architectural design in the case of microprocessors.

Also he could see a trend towards pure design companies, which actually got their chips manufactured at companies mastering a good process technology. He explained that this

evolution was likely to take place in the Silicon Valley, as all kinds of support companies and subcontractors for the semiconductor industry was located there.

We also discussed our TP2 project and our reasons for selecting Olivetti. He was very interested in the kind of applications to be performed on the system, and how the system solution was accepted within SDC.

This led to a long discussion concerning the evolving local network architectures and how our TP2 system could possibly migrate from its current centralized architecture to a pure distributed architecture. I noted, that we were aware of Olivetti´s plans concerning the SOLONE, and he suggested that I should go and see some of the companies involved in local network design in the Valley. He could however make no additions to the list of companies, with whom I alreadt had arranged visits.

I was then introduced to some of the persons in the list above, and was informed, that they would try as much as possible to answer my questions, although there might be some areas which would be considered too confidential. I noted, that SDC by our contract with Olivetti already had some sort of a non-disclosure agreement, but that it would be no problem to sign a special non-disclosure agreement for this visit, either personally or on behalf of SDC.

The Olivetti SW factory.

This group was managed by David Helms, with whom I had many talks on differnt software subjects.

They were currently looking into the UNIX operating system, and had plans of using it on their personal computers. They would provide a version of PLZ/SYS under UNIX. We had some talks on languages, and he told me that they were workning heavily with PASCAL/PLUS, a special Olivetti enhancement of Pascal. I never succeeded, however, to get hold of any manual on this.

The general trend of software for microcomputers, especially oriented towards small general applications system was the move towards UNIX as a standard operating system, not because UNIX might be the best operating system, but because it was readily available and easily transportable due to the large portion of it written in C. Also a lot of hobby computer manufactures and minicomputer companies were looking into UNIX.

By accident, I noticed a Zilog document on his table concerning parameter passing for the Z8000, and I was allowed a copy of it. This resulted in some long, tough discussion concerning parameter passing concepts, where Olivetti most likely would employ parameter passing in registers, against which I argued, that this would decrease the ability to write the low level assembler interface modules sufficient reliable and readable. I found out, that Olivetti did not use any kind of easy parameter definition and access in their current Z8000 development, and seemingly they did not consider programmer performance and reliability more important than pure processor performance. We had some argues about the need to optimize the use of a $100 processor chip, rather than optimize the use of a qualified programmer with a comparable hourly salary. Nevertheless, we agreed that it was important to have a precise and consistent definition of parameter passing, allowing procedures written in several different languages to call each other.

Olivetti had some joint-venture with Zilog concerning the

development of the PLZ/SYS language, but Zilog had chosen not to implement the floating-point real types used in Olivetti's PLZ. The fixed point real type requested by SDC to make PLZ look like PL/1 was an odd construct. Some of the work with PLZ had been done in cooperation with the people at Summit, but at the time for my visit, no one was actually doing anything on this.

The Olivetti DRAS group.

I was interested in a discussion with the DRAS group, as I knew they were carrying out a project on local networks. I managed to get the opportunity to talk with Lucas Cafiero, the manager of the DRAS group in Cupertino.

The DRAS group in Cupertino consists of 8 persons. One of their major tasks is to look at the technology and to investigate into specialized areas of interest to Olivetti.

One of their design tasks had been the development of the HPM bus performance evaluation system. Another task had been the design of a special computer module for the Z8000, using the Z8002 processor chip, but with some features meaning 20 percent savings in code and performance.

They were also responsible for the development of the SOLONE interface, a research project, which had been underway the last couple of years. They were also studying the Ethernet interface, and I saw examples of both. They were currently designing a terminal for the SOLONE, but would make no comments concerning the Ethernet.

I noticed some LSI mask drawings hanging on the wall, and by close questioning it was admitted, that a VLSI design project was currently in progress. They were investigating the cost/benefit of establishing a VLSI design group, which through the use of a CAD system, possibly from Applicon, should provide Olivetti with the ability to produce masks, which then could be transferred to actual chips at some of the semiconductor companies.

They would use the same approach to the semiconductor companies, as SDC uses towards Olivetti, i.e. being a large costumer they would be able to impact the semiconductor companies to produce their chips if Olivetti were to continue to buy from them. Intel was mentioned as an example of such a company, where this approach could be advantageous.

The applications of Olivetti's own LSI design would be special support circuits to be used in their computer

systems. A simple, but cost/effective memory management unit was mentioned as an example. The most important area would however be the design of properitary microprocessors to avoid the competition from other companies using standard microprocessors as Olivetti currently was doing, and to avoid compatibility between Olivetti software and software from other vendors. The most important area of use was expected to be in the typewriter and word processing business.

Olivetti were however reluctant to come with more specific information, as it had not yet been decided in Ivrea whether or not the initial experiments should be continued in a larger scale.

I was given a general overview of the DRAS organisation, so SDC was able to contact some of the persons of interest if needed:

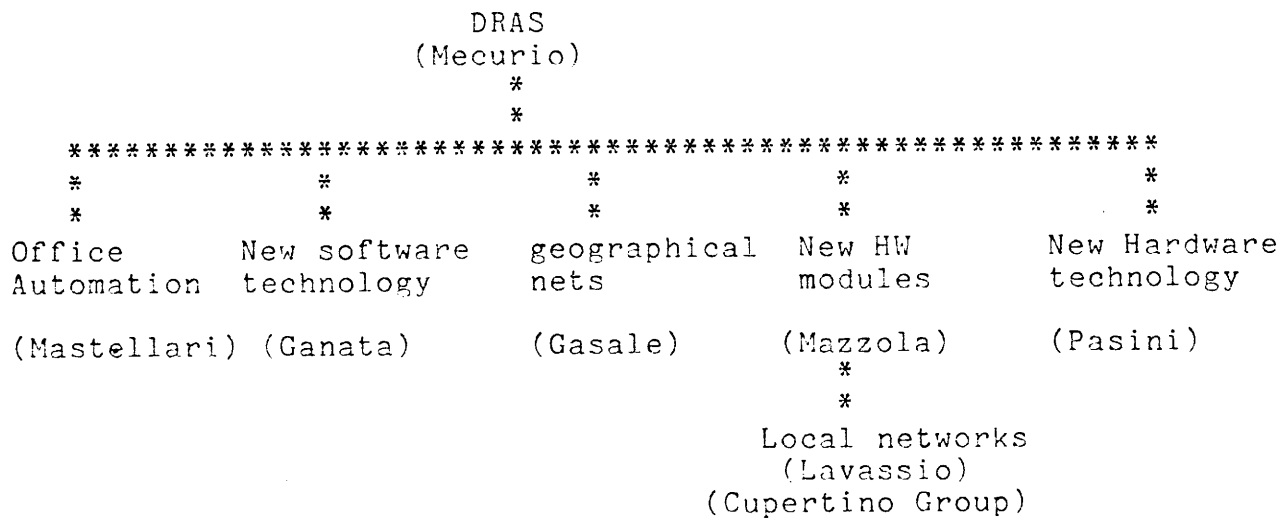

Meeting with Susan Owicki,
Computer Systems Laboratory, Stanford University.
****

During the spring and summer of 1980, I had without results tried to get some information on courses and contacts on Stanford University.

Luckily enough, I managed to get the name of Susan Owicki at the Computer Systems Laboratory from Jørgen Staunstrup at Århus University, and had received an acceptance of a visit before my departure.

Also, I had managed to get an appointment with Fouad Tobagi on the same institute, as I had seen a paper by him in some conference proceedings on local networks, and written to him with reference to this paper.

I was able to contact Susan Owicki by phone when I had arrived to the Bay area, and she mailed me a map of the university campus with directions of how to reach the institute, which proved to be very useful.

As expected, she could be of little help to my project as she worked on areas as theoretically verification of program correctness in concurrent systems. programming language semantics, and actually did not perform any kind of programming herself, nor being involved in microcomputer topics.

She gave me however an overview over the institute and the projects which could have interest for me, and provided me with a list of names, whom I could contact, including mr. Tobagi.

Especially, I was recommended to talk with Keith Lantz concerning distributed operating systems, with Jim Clark or John Hennessy concerning VLSI research, and with John Hennessy concerning programming languages. She also knew about a project called the SUN terminal, where the Z8000 was used to implement an intelligent terminal for distributed computing in connection with the Ethernet. This project was however done in the Computer Science Department by Vaugham Pratt, who was currently a visitor to Computer Science.

However, none of the persons were actually present on that day, and it turned out, that by the end of my tour, it had been neccessary to take to Stanford University on separate days in order to meet them. I never succeeded, however, to get in contact with Vaugham Pratt.

Meeting with Keith Lantz,
Computer Systems Laboratory, Stanford University.
****

Keith Lantz´s special area of interest was distributed operating systems.

I gave a presentation of my company, my specific project, our TP2 system and its architecture, and explained our views on how the technology might impact the future architectures of our branch office systems, and the problems of writing applications for future kinds of systems.

I asked about how they felt about the trends in computer technology and the future of software for distributed systems. I was told, that an Ethernet currently was being established for the entire Stanford University in order to both interconnect the numereous number of systems and terminals, and to investigate in this type of architecture, both with respect to hardware interfaces and to software design.

Keith Lantz had previously been at the Department of Computer Science at University of Rochester, where he had been involved in the design of RIG, an architecture for distributed systems. RIG was the implementation of a multiple-machine, multiple-network distributed system. RIG was intended to be the only intermediary between a human user (working through a terminal or a personal computer) and his available computer facilities. RIG was designed to present a coherent view of the distributed system similar to that provided by traditional operating systems for stand-alone computers. RIG was based on a model of distributed computation (independent processes communicating only by messages) which would allow the programmer to ignore the details of network and systems configuration.

I was given a copy of a paper, he was to present at the forthcoming ACM conference on distributed computing.

We talked about other problems of software concerning distributed systems and I was given some examples of interest, which he felt could be studied with good results. These were the Perseus System, and the Carnegie Archival

Storage System.

Of relevant references, he suggested

"The Design of a Network-based
Central File System"
Report CS319G, Aug, 1980
Bolt Bernnech & Newmann

and

Report £3796
"Network Operating Systems"
March 1978
Robert H. Thomas et al.
50 Moulton Street
Cambridge
Massachusett 02138

and

the April 1980 issue of IEEE communications

.

Meeting with John Hennessy,
Computer Systems Laboratory, Stanford University.
****

John Hennessy is responsible for the VLSI research at Stanford University. He is also involved in research on programming languages, operating systems and distributed computing.

I explained of my reason for visiting him, and our interest in the evolutions within VLSI technology, both seem from the semiconductor manufacturer's point of view and from the University standpoint, as the universities possibly might consider VLSI design in a way different from that of the manufacturers, who had a product to sell, and that the universities might be the first to consider new architectures better suited for VLSI implementation.

Hennessy felt, that there were two parts in VLSI design.

> First of all the design functions or design companies, who actually did the architectural and functional design, and preparing the mask layouts, and
>
> secondly the processing function or processing companies, who did the job of transferring the mask layouts to silicon, but not neccessary involved in the actual design and functional specifications.

These two functions had traditionally been in the same company until recently, but he felt that a trend was emerging, by which the functions would split up into dedicated design companies and dedicated processing companies, Some examples of this had already been seen in Silicon Valley.

He felt, that VLSI design was going to be easy and possible for average companies by using standard rules, e.g. the Lynn Conway method.

New design tools for VLSI design was underway, including functional oriented design. Also cheap design stations for

VLSI would become available, thus permitting small companies to invest in these stations and thereby being able to do VLSI design without having the need for todays very skilled experts.

We then discussed the VLSI activities at Stanford. They were able to do all the activities needed for design themselves, including CAD design, mask preparation and wafer processing. The turn-around time was one month for Stanford in-house developments.

The prime applications at Stanford of VLSI technology were innovative architectures (like graphics processors and highly parallel processors) and very high speed processors.

Also, they were in contact with some of the semiconductor companies further down in the Valley, and had done some contract work in connection with the new microprocessors. Their suggestions of improvements, however, had not generally been followed. Also, there was a great difference in the policies and structure of the various companies. Zilog, as an example, was relatively open, but had seemingly some internal problems of getting the ideas from Stanford accepted. Intel, on the contrary, was an extremely closed and unfriendly company. The laboratory was however able to get some insights into the developments within the companies, as most students were hired by these companies, of whom some still remained in contact with the people at the laboratory.

We then discussed the general trends of the computer market. and the views were very concrete:

- Distributed operating systems and local network architectures was the standard of the future. Especially Ethernet was thought to have good market possibilities, as the concept was fairly clear and consistent. The laboratory had in fact designed their own Ethernet/CSMA chip.

- High technology compilers for portable programs were needed, and would probably bee more common in the years to come.

- Pascal contained some good concepts, but were useless

for actual programming. ADA might be too comprehensive, but a superset of Pascal was definitely needed.

The concept of moving software into silicon as an integral part of the microprocessor chip was somehow doubtful, as this would mean severe problems with software reliability and updating, and it might cause a costumer to be locked into a specific manufacturer. The general view of Intel, who has first promoted this concept, was that Intel was too dominated by hardware design and they might not fully have realized the impact this approach would make on their software support responsibilities.

We then discussed the possibility of getting a continous update of the research activities at the laboratory, and I was told that generally outside companies would have to pay for any kind of consultation due to the way the university and the laboratories were financed. Also much of the project work, which was carried out, was done by contract with other companies, and thus underlaid non-disclosure agreements.

One way, however, to get an update on the technological work within the computer field, was to become a member of the Computer Forum, a kind of club established to give interested companies a competent follow-up. This would mean copies of various research reports, lists of graduating students and their qualifications, and a visit once a years by a professor or lecturer from the laboratory on the company's location where he would present talks and papers on matters of interest for that particular company. The annual fee of the Computer Forum was 10-11000 dollars, but he was not quite sure how this was arranged with European companies. The only European company of the Computer Forum until now was Siemens in Germany.

I received some technical reports from him concerning his work within the VLSI and programming language areas.

Meeting with Fouad A. Tobagi,
Computer Systems Laboratory, Stanford University.
****

Tobagi is assistant professor of electrical engineering, and has done some theoretically work on the performance of local networks using the CSMA contention scheme. His special areas of interest are computer networks and communications, modelling and analysis of computer systems.

We had a long talk, where I presented SDC, the TP2 project and my project. He was very impressed to meet one from the "Danish Bank", as he actually had done some contract work for the Olivetti technology center in Cupertino, and through them heard about the big project. He found it important, and a good idea, that SDC as a costumer would investigate in the technology and thereby being able to impact Olivetti.

He felt, that the concept of local networks in the form of the multiple access, broadcast bus was the only one currently suited for business applications, although he was not quite sure whether the Ethernet or the IEEE standard would be chosen. At Stanford, they were implementing the Ethernet as a research project, and he would personally recommend users to employ the Ethernet, if they were to implement a distributed computing system today, because the Ethernet concept was sufficient consistent and advanced for today's needs. In long term planning, users should consider both the Ethernet and the IEEE standard, but it would be waste of time to try to implement something else not in accordance with these concepts.

He recommended me to visit the Ungermann/Bass Corporation, Zilog and possibly AMD to discuss local networks in a practical sense.

He had participated in the IEEE local network task, but had kept a rather low profile as he believed it more important to get an official standard as soon as possible, rather than getting an theoretical optimal, but unimplementable one. He had however done some work on an enhanced concept based on token communication, which he considered superiour to the Ethernet way of communicating, and intended to present this work at one of the future meetings in the committee.

We discussed the various problems of writting software for distributed systems, and our specific problems of transporting software from the centralized architecture in our TP2-system to a true distributed architecture. He agreed, that it was neccessary already now to design software for a distributed architecture, making it sufficient modular and structured to be transferable at a future point in time.

He told me, that Olivetti actually had been considering local networks for years, although he was not allowed to disclose this further, but he was surprised to hear that SDC had not been informed about this by Olivetti, especially seen in relation to the TP2-project which was an obvious target for local networks.

Meeting at SRI International

xxx

****

<to be written>

Meeting at Motorola

xxx

****

<to be written>

Meeting at AMD
xxx
****

<to be written>

Meeting at Prolog

xxx

****

<to be written>

Meeting at National Semiconductor

xxx

****

<to be written>

Meeting at Interface West
xxx
****

<to be written>