

## **Sabbatical in Japan: Collected Trip Reports**

*Richard D. Schlichting*

TR 91-22

### *ABSTRACT*

During the course of a sabbatical in Japan in 1990, the author visited 20 universities, industrial research laboratories, and government laboratories. This report collects together trip reports from a number of those visits. Statements made herein represent the personal opinion of the author.

October 6, 1991

Department of Computer Science

The University of Arizona

Tucson, Arizona 85721

## Sabbatical in Japan: Collected Trip Reports

### Fujitsu Laboratories (March 29, 1990)

On March 29, I visited the Fujitsu laboratories in Kawasaki outside of Tokyo. My official host for the visit was Dr. Takao Uehara, who is the manager of their software laboratory. Sanya Uehara, who works in the software lab, was also heavily involved in my visit, and the one who did most of the arrangements. He graduated from Tokyo Institute of Technology about 10 years ago and has worked at Fujitsu ever since. Also involved was Moto Adachi from the human interface research division who I had met in Tucson prior to leaving for Japan in December. (See addresses below.)

I left the office about 8:15 and arrived at the front gate of Fujitsu around 9:00 without too many problems. After getting a visitors badge, I was pointed to the main building -- a gleaming 20-story building that is apparently about 2 years old. At that point, I was directed to a conference room on the second floor where most of our discussions were to take place. I was allowed to venture up there on my own, which surprised me somewhat after the "escort" requirements that seem standard at many U.S. laboratories.

At about 9:25, Sanya Uehara came in and introduced himself. We chatted for awhile, and he gave me glossies describing Fujitsu, Fujitsu Labs (which is technically a wholly-owned subsidiary), etc. At 9:30, the official part of the visit began as Takao Uehara arrived. He gave me a short description of the lab setup and then put on a videotape describing some of the activities.

While the tape was on, other people started arriving, including Moto Adachi, Mijio Aoyama from the business switching systems division, and Yasunori Kimura from the Artificial Intelligence lab. Probably a total of 8-10 in all. After the tape and introductions, I gave my talk on failure-handling mechanisms for the SR distributed programming language.

Afterwards, there were two presentations about projects going on at Fujitsu, one by Aoyama on their multi-processor digital switching systems and simulator, and another by Kimura on their 5th generation project to construct a parallel machine. In the talk on switching system, Aoyama first laid out the history and general hardware of such systems, both at Fujitsu and elsewhere. He then described in more detail their systems, including an overview of architecture of the software system written in Chill. Their most advanced model uses multiple call processing processors as a way to incrementally increase capacity and duplex hardware for reliability. Finally, he described a special-purpose simulator for the system that they had constructed; it runs on multiple Suns (currently 2) in a local-area network. Interestingly, it does not simulate failures, so the reliability aspects have not been tested using this vehicle.

Kimura then talked about the projects in the AI lab. He first described the KL family of languages that are being developed at ICOT as part of the well-known "5th generation" project sponsored by MITI. The languages are based on Guarded Horn Clauses (GHC) and have a resemblance to Prolog. The family consists of 3 or so languages, from the base one to a "user version" that runs at the top level. I should note that his knowledge of this work was first hand, as he had been assigned to ICOT for 5 years as part of Fujitsu's contribution to the project.

He then went on to describe more of the hardware architecture, which is based on multiple processing elements organized in clusters. They are currently in the process of designing the PE's, which are intended to be optimized for KL. They are apparently under a great deal of pressure to get the entire machine built before the end of the project, which is now only 2 years

away. In addition to the usual difficulties, they also have the problem that the language specs keep getting changed. (When asked when he thought that would stop, he said in 2 years when the project ended!) There are 5 companies in competition to develop the hardware.

After eating lunch in a nice restaurant on the top floor of the building, we spent a few minutes in an exhibition room in the basement. The bulk of it consisted of "hands-on" exhibits of a lot of the current Fujitsu products. In addition was my personal favorite--a working model of a 1954 Facom computer built using mechanical relays! They cranked it up for me and after a lot of clattering, it printed out 3 or 4 powers of 6.

After lunch, S. Uehara escorted me to another building to rejoin Takao Uehara, who was to accompany me to three demos. The building in which they were located is apparently part of the software lab managed by Uehara, and was quite a bit older than the main building. The first demo was a parallel machine called CAP (Cellular Array Processor). After describing the architecture, I was shown some sample graphical and heat simulation demos. The existing machine is apparently the 2nd generation, of which 1 has been shipped. They are currently starting work on the next generation.

The second demonstration was of a simulation of "neuro-computer." After a short description of the concepts, I should shown an application consisting of a curve-fitting application using a 3-level neural network on-line. I was then shown a video tape that elaborated on the concepts, described out they are used to build "machine that can learn", and showed a cute robotic application. One was dressed up as a robber, one as Sherlock Holmes, and two as policemen, with the object being for Holmes and the policemen to capture the robber. The robber had been taught to flee, while the others had been taught to follow the robber with varying degrees of skill. At the very least, it certainly was entertaining.

The final demo was of a machine translation system (Japanese --> English) that was recently released commercially. According to the person doing the demo, it's being used to do such things as translate handbooks and manuals, applications where some of the resulting awkwardness can be tolerated. He described the basic structure of the system, which uses an intermediate form called "interlingua", and then ran some sample sentences. The samples were intended to show how the system can make some fairly fine distinctions, and I actually found it pretty impressive despite the fact that the person doing the demonstration downplayed its sophistication. The algorithms were developed by the people in the lab, who then gave it to the products division who did things like add an editor front-end to it. They are apparently working on a version that will translate English to Japanese based on the same interlingua idea.

That ended the official part of my visit. I then sat with Moto Adachi for an hour or so chatting about various things. Among the more interesting were his comments on how it was getting hard to hire new people because starting salaries are low compared to places like banks and insurance companies. This aspect of the Japanese labor shortage has been given much attention in the English-language press as well. I found this to be an interesting difference from the situation in the U.S., where my perception is that it would be unusual for a person with a technical degree—especially a graduate degree—to consider going that route.

Addresses:

Dr. Takao Uehara  
Deputy General Manager, Information Processing Division  
Manager, Software Laboratory  
Fujitsu Laboratories, Ltd.  
1015 Kaminodanaka, Nakahara-ku  
Kawasaki 211, Japan

Sanya Uehara  
Software Laboratory  
Fujitsu Laboratories, Ltd.  
1015 Kaminodanaka, Nakahara-ku  
Kawasaki 211, Japan

Motomitsu Adachi  
Senior Engineer, Human Interface Research Division  
Fujitsu Laboratories, Ltd.  
1015 Kaminodanaka, Nakahara-ku  
Kawasaki 211, Japan

Mikio Aoyama  
Senior Engineer, Business Switching Software Dept.  
Business Switching Systems Division  
Fujitsu, Ltd., Fujitsu Nakahara Bldg.  
629 Shimokodanaka, Nakahara-ku  
Kawasaki 211, Japan

Yasunori Kimura  
Artificial Intelligence Laboratory  
Fujitsu Laboratories, Ltd.  
1015 Kaminodanaka, Nakahara-ku  
Kawasaki 211, Japan

### **Fujitsu Numazu Complex (April 25, 1990)**

On April 25, I visited Fujitsu's Numazu complex, which is at the base of Mt. Fuji near the city of Mishima. The trip was arranged by Prof. Yoshihiro Tohma of the Dept. of Computer Science at Tokyo Institute of Technology where I was staying on sabbatical. He had been invited by Fujitsu to present a lecture on fault-tolerant computing at the plant, and came up with the idea of having me accompany him to present a complementary lecture on fault-tolerant software techniques.

We started out from the institute at about 8:55 AM and headed to Shin-Yokohama station. There, we met Yohichi Suzuki, a former student of Yoshi's (the first, it later came out) who was to be our host for the visit; his current position is that of deputy general manager of Fujitsu's mainframe division. After introductions, we caught the 10:05AM Kodama shinkansen for Mishima station, where we arrived at approximately 10:50. Another 20 minutes by taxi brought us to the Numazu complex, which is situated in a wonderful location not far from the coast. We were escorted to a meeting room, where we settled into a short discussion with Suzuki, Masahito Date, a software manager who was familiar with my work on fault-tolerant computing, and

Makoto Kimura, a hardware specialist who is also a graduate of Tokyo Institute of Technology.

The talks began at 1:15PM, and we got to the room right about on time. We had been informed during lunch that there were going to be about 180 people at the presentations, and they were all there when we arrived. After introducing us, Yoshi began his talk (in Japanese, of course.) His topic was a general introduction to fault-tolerance, during which he gave motivations, overviewed some current companies and architectures, etc. Afterwards, he showed a 20 minute videotape of the SAFE system that they had worked on at TIT.

After a short break, it was my turn. I gave a general overview of fault-tolerant software, which included talking about various useful abstractions. It went quite well, I think, and I later got good feedback on the talk over dinner. There were a couple of good questions after the presentation, which were given in Japanese and translated for me by a fellow in the front row.

Following the talks, we were shown around the plant, with a tour guide leading the way. By that time, we had been joined by Katuhei Oshio, who is manager of the TMP Development Dept. of the Product Assurance Division. The Numazu complex is the manufacturing site for the high-end mainframes and supercomputers, and we got to see how they were assembled. It was fascinating, especially for a hardware novice like me. I thought it especially interesting that most of the assembly for such complex machines is done by hand, with limited automation only at a very few stages. We were also shown the rooms in which they were tested, the main software labs and system test rooms, a little history museum with an old relay computer, and a piece of natural language translation software called ATLAS II.

We left around 8:30 by taxi and made it to Mishima station just as the train was pulling up. In addition to Mr. Suzuki, we were also joined by Mr. Oshio since he lives in Yokohama City. We talked about a number of general things on the way back, including plans for the Pacific Rim Fault-Tolerant Computing conference to be held the following year. (Mr. Oshio is the local arrangements chair.) We arrived at Shin-Yokohama station around 9:30, where we caught a taxi back to Ookayama.

Addresses:

Yohichi Suzuki  
Deputy General Manager, Main Frame Division  
Fujitsu Limited  
1015 Kamikodanaka, Nakahara-ku  
Kawasaki 211, Japan

Masahiro Date  
Manager, 3rd Section  
Development Planning Dept. I  
Software Division, Computer Systems Group  
Fujitsu Limited  
140 Miyamoto, Numazu-shi  
Shizuoka 410-03, Japan

Katsuhei Oshio  
Manager, TMP Development Dept.  
Product Assurance Division  
Fujitsu Limited  
1405 Ohmaru, Inagi  
Tokyo 206, Japan

Matkoto Kimura  
Computer Engineering Dept. IV  
Main Frames Division  
Fujitsu Limited  
1015 Kamikodanaka, Nakahara-ku  
Kawasaki 211, Japan

Prof. Yoshihiro Tohma  
Department of Computer Science  
Tokyo Institute of Technology  
2-12-1 Ookayama, Meguro-ku  
Tokyo 152, Japan

### **Hokkaido University (May 10, 1990)**

On Thursday morning, I was met at my hotel by my host, Prof. Eichii Miyamoto from Hokkaido University. Since the hotel was very close to campus, we walked to his office, which gave me a nice chance to look at the university and get some general impressions. Overall, I would say that the feeling is one of spaciousness and newness, making it appear much more like an American university than many of the other Japanese universities I visited. This stems partially from the fact that an American was instrumental in setting up the university in the early part of this century and partially because Sapporo itself is much newer and more spacious than other Japanese cities. The university seems very comprehensive, as I saw everything from an agriculture school to a medical and dental college.

Our first stop was the main College of Engineering building, where we stopped to get an English-language brochure on the College. We then walked up and over a covered walkway to the Information Engineering Building, a 1 and a half year old building where Miyamoto's laboratory is located. Once there, we sat and chatted for awhile about the general setup of the lab. A few points that he made:

- At the current time there are approximately 7 undergraduates, 7 MS students and 3 PhD candidates associated with his lab. He indicated that this number was about average for the other labs in the department as well.
- They have been connected via email to other Japanese universities for about 3 years. This is used extensively to keep in touch with colleagues elsewhere in Japan, and is especially important given Sapporo's relatively long distances from other cities.
- Japanese-language interfaces are no longer the stumbling block that they were just a few years ago.

I ended up giving him the rundown of my university and department as well, and we both found the differences between the two educational systems interesting.

After about 45 minutes or so, we left to go talk with his students and the other professors associated with his lab. They each introduced themselves and briefly characterized their research area; most were working in object-oriented systems, with some AI as well. Two students then got up and took turns describing their work on a system they called Kamui. Kamui, it turns out, is an object-oriented system in which objects can be written in different programming languages. The interface of each object is defined using a protocol that described the operations the object exports and the types of the arguments. Another object wishing to use these exported operations name the same protocol. A single object can also have multiple protocols describing different interfaces; although similar to inheritance, it differs in that it need not form a hierarchy. Kamui also has a shell, which is apparently used to make Kamui objects out of non-Kamui programs. This shell also interacts with some user interface work they have done to allow, for example, the mouse on one machine to control the cursor in the window of another. Differences in value representations between different languages is handled by translating into the target representation before sending.

This sounded very related to some of my own work, so I then got up and gave an impromptu overview of our MLP system. They had not heard of the system previously, but were familiar with better know related systems such as the HCS work at Washington, Mercury at MIT, etc. In general, they seemed very much aware of the issues involved and had, in my opinion, constructed in Kamui a very interesting solution.

Another student then got up and described his work on a hierarchical nameserver scheme for wide-area networks. The main selling point appears to be that requests for application programs that happen to be on failed nodes are automatically rerouted through the nameserver to another copy. The focus of the scheme is primarily on finding relatively static objects such as code, so that they have not had to address the problem of keeping the copies of the objects consistent in the face of failures.

At that point, the meeting broke for lunch. Afterwards, I gave a couple of formal research presentations, and was then treated to a nice tour of Sapporo by my host.

Addresses:

Prof. Eiichi Miyamoto  
Dept. of Information Engineering  
Faculty of Engineering, Hokkaido University  
Sapporo 060, Japan

### **Osaka University (May 16, 1990)**

I arrived at the Toyonaka campus of Osaka University at about 7:00p on May 15. As soon as I got out of the taxi at the main gate the guards knew who I was, as they had obviously been forewarned to look out for a foreign visitor. They then phoned my host, Prof. Tohru Kikuno, who came to pick me up and take me to my on-campus accommodations for the night.

The next morning, Kikuno met me at my room and we headed out to his laboratory, which is headed by Prof. Koji Torii, a quite well-known expert in the area of software engineering. There I met Torii and we all chatted for awhile about Japan and computer science in general.

Following our conversations, he took me into another room to hear a couple of presentations. The first, by a research associate in the lab, was on the Ginger data collection and analysis system, which is now a couple of years old. The thrust of this system is to collect

information on code size and changes throughout the development of a program, and then to use this data to provide feedback to programmers during development. A prototype system has been written in C, and its usefulness tested during an experiment involving various teams of university students. The results seem to support their hypothesis of the usefulness of this information. Interestingly enough, subsequent internal use of Ginger has been only to collect data, and not to provide feedback.

The second presentation was on Square, which is a system that attempts to use heuristics to develop project development plans similar to those developed by human managers. They have experimented with two heuristics, with the second one giving results similar to those obtained from the use of a human project manager. The overall goal, which is in keeping with Torii's general research thrust, is to improve the general level of software quality by investigation of high-level design processes.

Following lunch and my talk on Psync, we adjourned again to Torii's office for general discussion. While we were talking, Prof. Kozo Kinoshita, who I have known for a number of years because of his work in fault-tolerance, stopped by to chat. He moved from Hiroshima University to Osaka Univ. about a year ago and is now in the Department of Applied Physics. Moving universities is relatively unusual in Japanese academia, but this particular move was not since Koza received his degree Osaka University. In fact, he and Torii actually graduated from the same lab at Osaka University, with Kozo being a couple of years ahead.

At about 5pm or so, it came time to head to Kyoto for my next day's appointment. Dr. Kazuo Sugihara, a former student of Kikuno's who is now a faculty member at the University of Hawaii, was going to Nagoya to visit friends, so we decided to travel together on the Shinkansen as far as Kyoto (we both got a chuckle out of the fact that he--the Japanese--was traveling on a Japan Rail Pass, whereas I--the foreigner--was ineligible due to my visa status.) We had an interesting talk about the differences between the U.S. and Japanese educational systems on the way. For example, he tended to agree with the common assessment that Japanese undergraduates do not work as hard as American undergraduates, although he did indicate the situation was more comparable at the graduate school level.

Addresses:

Prof. Koji Torii  
Dept. of Information and Computer Sciences  
Faculty of Engineering Science, Osaka University  
Machikaneyama 1-1, Toyonaka City  
Osaka 560, Japan

Prof. Tohru Kikuno  
Dept. of Information and Computer Sciences  
Faculty of Engineering Science, Osaka University  
Machikaneyama 1-1, Toyonaka City  
Osaka 560, Japan

### **Nagoya University (May 18-19, 1990)**

I arrived at Nagoya Station from Kyoto a little after 9:00pm and immediately took a taxi to my hotel, which turned out to be a place that catered primarily used by traveling government employees. I had a room on the top floor with a wonderful view of the famous Nagoya castle,



which was a quite spectacularly illuminated by floodlights.

The next morning I was picked up by one of my hosts, Prof. Toshiku Sakabe, who is an associate professor in the laboratory of Prof. Yasuyoshi Inagaki. (Prof. Inagaki was out of the country at the time.) We drove the 20 minutes or so to the campus on the other side of town, making conversation as we went. Once we arrived, we went directly to his office in one of the engineering buildings. My other co-host was Prof. Tomio Hirata, who is also an associate professor in the Inagaki lab. Interestingly enough, however, his appointment is actually in a different department (Information Engineering) than Sakabe's (Electrical Engineering.) They commented that this is actually not that unusual, and that one problem with Japanese computer science is that the people doing this type of research can be scattered throughout a university.

Fairly soon after arriving, we went into the conference room for an overview of the work underway in the lab. The basic flavor of research is theoretical, and concentrated primarily on various topics related to formal semantics. Three presentations were given on the work. The first was on algebraic semantics and the complexity of term rewriting systems, the second on modeling monitors using Milner's CCS, and the third on extensions to non-monotonic logic to solve the so-called "multiple extension problem." All three were interesting, although because of my background, I was really only able to follow the CCS work in any detail. In addition to myself and Sakabe, the talks were attended by a goodly number of students.

Afterwards, I got into a very wide-ranging discussion with Sakabe and his colleagues about the structural and funding differences between U.S. and Japanese universities. For example, a typical yearly budget from their university for one lab is on the order of 6 or 7 million yen (~\$40,000 - \$50,000), with a grant from the Ministry of Education adding about an additional 5 million yen (~\$35,000) per year. Out of this comes travel, supplies, etc., but not salaries. We also talked some about the "chair system", called the koza system in Japanese, and about the equivalences of various ranks. According to Sakabe, the standard progression is joshu (5 years), jokyoku (10 years), and then kyoju until retirement at 60, 63 or 65, depending on the university. The latter two are typically translated as Associate Professor and Professor, respectively, although it's not clear that those ranks are equivalent to the American ranks, especially since there is an intermediate rank, junkyoju, that is sometimes used as well. After some discussion, we left it with the speculation that the most accurate equivalences might be the following:

joshu == postdoctoral fellow  
jokyoku == assistant professor  
junkyoju == associate professor  
kyoju == professor

We also talked some about the way departmental financing is done in the U.S. and about the grant system. One thing that amazed them was the high percentage of grant funds that are taken out by the university for overhead.

After lunch, I was given a tour of their lab's computational facilities. There were a number of Suns of various varieties and configurations (including some OEMed by Fujitsu, Toshiba, and Omcron), some Sony News machines, some Macs and a few PC-class machines. Some of the machines are in individual offices (for example, Sakabe had a Sun and a PC-class machine), with the others in shared facilities. Nationwide network access is somewhat limited at the current time, although there are plans to upgrade soon.

Following a nice dinner, I spent the night at my hotel, returning to Tokyo the next morning.

Addresses:

Dr. Toshiki Sakabe  
Associate Professor  
Dept. of Electrical Engineering  
Faculty of Engineering, Nagoya University  
Furo-cho, Chikusa-ku  
Nagoya 464, Japan

Dr. Tomio Hirata  
Associate Professor  
Dept. of Information Engineering  
Faculty of Engineering, Nagoya University  
Furo-cho, Chikusa-ku  
Nagoya 464, Japan

### **NTT Software Laboratories (May 22, 1990)**

On Tuesday, I traveled to the west side of Tokyo to visit NTT Software Laboratories. I arrived at the closest station -- Mitaka on the Chuo line -- at about 9:30am and from there took a taxi to the labs.

Once I arrived, the guard phoned my host, Dr. Naohisa Takahashi, who came to sign me in. The lab itself consists of about 10 relatively new buildings in a pleasant setting. Takahashi led me to his office, which is in one of the main buildings. His background is mainly in functional languages, but recently he has been a member of a distributed computing group and working on replay-style debugging of concurrent, functional programs.

For about 1/2 hour or so, he introduced the labs to me, overviewing the structure and groups. He then took me into a lab with (mainly) Suns to meet two other researchers. One, Ken-Ichiro Murakami, first described NTT's internal network to me. The system is built on TCP/IP and is consciously modeled about Stanford's, from which they apparently received guidance. Among the interesting points made:

- Compared to the U.S., there is a general lack of network expertise in Japan.
- A fair amount of the network equipment they use is American (e.g., Cisco routers) since they have been used more and hence have been debugged. Good technical support is also an issue.
- They consider network connectivity to the U.S. as an important part of general computing infrastructure. They were the first Japanese affiliate of CSNet, to which they are connected using two separate commercial transport services.
- Their efforts at expanding their international computer networking is hampered somewhat by regulations put forth by the Ministry of Telecommunications and Post. The basic problem is that they are not authorized as an international carrier, and so are prohibited from transferring outside (i.e., non-NTT email) across the Pacific over their leased lines. They expect some regulatory relief in the next few years.

Following that presentation, another engineer named Toshiharu Sugawara demonstrated a heuristic-based Cooperative LAN Diagnostic Expert System. (A paper on this appeared in the

IEEE Phoenix Conf. on Computers and Communication last year.) The basic idea is that a person having network troubles (e.g., unable to connect) would use the program to determine the problem. There is one instance of the program running on each network segment and the pieces cooperate as needed to determine the problem. The basic rationale for the development of the program is the lack of enough trained technical people who can be relied on to manually diagnose the problem. My understanding is that the program is still in development and is not yet used on a daily basis.

Next, we went to lunch at the Tsuken Club located in a different building. The building appeared to house a general cafeteria, but we ate in a private dining room. Joining us for lunch was Prof. Sam Toueg from Cornell, who had arrived in Tokyo the previous Friday to start a month long visit to NTT. Apparently, an NTT researcher had visited Cornell for a time the previous year, and so they in turn had invited Sam. It was of course nice to see Sam again, and the entire group of us had a quite enjoyable conversation over a very excellent continental-style lunch.

My presentation on Psync came after lunch, followed by a presentation and demo of a software development environment for communication software by a Dr. Haruhisa Ichikawa. The basic goal of the work is to perform automatic generation of communication software from specs, especially in an incremental fashion as the specs evolve over time. The system includes a graphical display of the specs, which were written in a standard protocol description language. The system also detects problems with updated specs such as inconsistencies and the like. The system runs on Sun workstations and is being done jointly with British Telecom. Although I am not an expert in this area, I found the demo fairly impressive. In response to a question, however, Dr. Ichikawa seemed rather skeptical that this would be adopted to develop real-life communications software at any point in the near term; he essentially said that development departments have their own established methods that would be very difficult to change.

The final work described during my visit was a debugger for parallel and distributed programs. This research was being directed by my host Dr. Takashi, and the presenter was a young research engineer who was collaborating on the work named Yoshifumi Manabe. The system is based on a replay technique and the use of global predicates to establish breakpoints. Among other things, they have perused what appears to be some fairly theoretical work on what type of global predicates are feasible for use in such debuggers. He gave me a couple of papers (including one in English) that describe the work.

Overall, the visit to NTT Software Labs was very enjoyable. My impression from their presentations is that all of the various projects are technically sound, and that they understand well the basic issues and background. I also believe that they grasped well and appreciated the material on Psync that I presented; this includes the fault-tolerance aspects, which is especially impressive since, as far as I could tell, they are not doing any work specifically in that area. I also got the feeling that NTT encourages international cooperation and publication (their visiting researcher program is also evidence of this.)

[Addendum added June 6, 1991] The LAN diagnosis system mentioned above is now being used occasionally to fix network troubles in the Lab, and has proved to be a useful tool. Dr. Takahashi's current research interests are focusing on the debugging of parallel programs with shared data objects.

Addresses:

Dr. Naohisa Takahashi  
Software Research Laboratory  
NTT Software Laboratories  
9-11 Midori-cho, 3-Chome  
Musashino-shi, Tokyo 180 Japan

### **Kyushu University (May 28, 1990)**

On May 28, I visited Kyushu University, which is located in the southern Japanese city of Fukuoka. I was picked up by my host, Prof. Kazuo Ushijima, at about 9:30 AM. We then walked the 10 minutes required to reach his office, entering the campus through what appeared to be a small south gate. Upon arriving, we chatted for bit and then he briefly showed me a PC-based Japanese writer's assistant. The first version was written in a version of Snobol4 that they had augmented with Japanese features back in 1983 or so, although the newer version he demoed is written in C. The former version was of special interest to me since one of the developers of Snobol4 -- Ralph Griswold -- is a faculty colleague at the University of Arizona.

After a few minutes, Ushijima passed me on to the Associate Professor in his laboratory, Prof. Keijiro Araki. After a few of the usual pleasantries, I asked about equipment, and he ended up giving me a very complete rundown and tour of the facilities. They are, in short, impressive. The bulk of their equipment is from DEC, and they are apparently one of DEC-Japan's most important academic customers (the president of DEC-Japan was going to visit a couple of days after I left.) The centerpiece of their equipment is a Vax 8800, which is used for both departmental research and teaching. This is, in fact, the first of this type/class machine I have seen on my visits; they acknowledged that this is unusual, although claimed that it was common for departments to have 780-class machines several years ago. They also have 30 Vaxstation 2000's that are used primarily by undergraduates, mostly to get windows on the 8800 (seems that their lack of memory--only 4 MB--makes them very difficult to use as real workstations.) Another 11 Vaxstation 2000s were distributed around to different laboratories for use by graduate students. There are also 3 or so Vaxstation 3500s, as well as some Xerox machines for word processing. In Ushijima's laboratory specifically, I also saw 7 or 8 Suns, as well as a scattering of Macs, PC-class machines, etc.

In addition to a tour of the facilities, Araki also gave me an overview of the personnel in the lab, some information on the department, and also a feel for how the department fit into the University. Among the interesting points:

- The department is called the Dept. of Computer Science and Communication Engineering, having evolved from the Communication Engineering department.
- The department historically has had 7 chairs, with 2 being devoted to CS (Prof. Ushijima's chair is called Computer Software.) However, they now have 11 chairs; two of these chairs were from other places in the University, while they received one each from the Electricity Dept. and the Electronics Dept. (two departments with which they closely interact.) The new chairs are partly a reorganization--bringing together AI-related work--but also some growth.
- In the lab currently are 3 PhD students, 7 MS students, and 10 undergrad. This is apparently a bit lower than normal as he mentioned that they had more students last year. He also said that this was about average for the other labs in the department.

We were also able before lunch to talk some about his own work, which focuses primarily on formal semantics of concurrent and distributed programs. Despite his use of formal approaches, his goal--like mine--is to develop pragmatic techniques. One thing that he has done is to attempt to apply some different techniques to a robotics system. The work is being done jointly with with people from Yasukawa Electric works, which is located in Kita-Kyushu; this company originated as a manufacturer of heavy machinery for steelbuilding, chemical plants, etc, but has recently moved into the manufacture of industrial robots.

The example program itself is a robotics application in which they wish to establish properties about the movement of the robot. For example, they want to argue that synchronization is done correctly so that the arms never touch each other. The first attempt was done using temporal logic; the basic logic was from Kroger's book, although they found it necessary to do extensions. The second was using the algebraic approach developed by Manfred Broy, a natural since Araki had recently returned from spending 10 months on a Ministry of Education grant at Passau, West Germany studying with Broy. He gave me several papers describing this work, mostly in Japanese.

At this point, Araki introduced me to Dr. Norihiko Yoshida, a research associate working in distributed systems, and Zengo Furukawa, a lecturer associated with the group, but technically from another department. We talked about research, with Yoshida and Furukawa giving me an overview of their work. Given my background, Yoshida's work was more interesting to me. Basically, what he is trying to do is implement distributed virtual shared memory for an AI-type blackboard application. The specific novelty of the work according to him is that the implementation is based on point-to-point "request" messages rather than a broadcast algorithm; the rationale is that this is supposed to work also in computing environments that are not based on a LAN. That claim seems reasonable, although it was not really clear to me exactly how his work relates to research in the area of reliable broadcasting or other work in DVSM. He gave me a paper in English describing the work.

Following lunch and my research talk, we adjourned to Araki's office for a discussion among Araki, myself, and Dr. Jingde Cheng, the other research associate in the lab. He was been working on several things, including an event monitor and debugger for Ada programs, formal methods for detecting deadlocks in Ada programs, and on non-monotonic logics. The work on the event monitoring was especially interesting to me since its use of partial ordering of events makes it very much related our own work on the Psync IPC mechanism. We also chatted for awhile about problems with Ada before I had to leave to catch the Shinkansen to Hiroshima.

Addresses:

Prof. Kazuo Ushijima  
Dept. of Computer Science and Communication Engineering  
Kyushu University  
Hakozaki, Higashi-ku  
Fukuoka 812, Japan

Prof. Keiji Araki  
Dept. of Computer Science and Communication Engineering  
Kyushu University  
Hakozaki, Higashi-ku  
Fukuoka 812, Japan

### **Hiroshima University (May 29, 1990)**

On Tuesday, I woke early and ate breakfast at the hotel coffee shop. It was a buffet spread, in fact, the same things as I remember from my previous trip to Hiroshima 2 years before. I then packed up, checked out, and waited for the arrival of my host, Prof. Tadao Ichikawa, a professor at Hiroshima University whose specialty is visual programming. The schedule was for him to pick me up at 9am for a visit to the famous shrine at Miyajima, and then to go to the university in the afternoon. He arrived a little after 9am, but informed me that, unfortunately, he was leaving for the U.S. in a couple of days and, quite understandably, couldn't really spare the time for the morning sightseeing. I was happy to go by myself, but it turned out that he had arranged for his assistant Mrs. Tanimoto to take me instead.

Following sightseeing and lunch, we drove to Ichikawa's lab at Hiroshima University, a journey that took a little over an hour. The particular campus where his lab is located is a new one that is situated quite far (35 kilometers) from Hiroshima proper. I found out later that there also have another campus downtown, but that all departments are scheduled to move, with the new buildings actually being financed by the sale of the old ones in the city. The actual building where Ichikawa's office is located is about 8 years old.

We were a bit late for my presentation, so Ichikawa postponed it for 15 minutes to allow me to rest. After my talk, we went to Ichikawa's office to chat for awhile. It turned out to be a very interesting time, since he has a number of very interesting opinions about Japan and the Japanese educational system. Indeed, he has written a book of non-academic essays called VIVA Nippon!? containing observations on different aspects of Japanese life (it was translated from a Japanese version published a couple of years ago.) I only glanced at it, but it looked very interesting (and entertaining), so I made a mental note to get a copy. He also gave me a copy of a report based on a talk made at the annual ACM conference this year that criticized the Japanese educational system (in a hilarious way.) He also made the following points in our discussion:

- Running conferences is very expensive in Japan compared to other countries since, for example, the hotels charge quite a lot to rent conference rooms, etc. This makes getting corporate support mandatory, although they are reluctant to contribute financially unless the meeting is sponsored by a purely Japanese professional organization.
- Many Japanese computer science results don't make it into international journals. In his opinion, the primary reason for this is that there are essentially no rewards built into the Japanese professional system to encourage researchers to be international in outlook. This, combined with the fact that it's more difficult to prepare a paper for this journal, means that most results get published only in Japanese journals or conferences.

In between various conversations, he had his associates demonstrate a couple of the visual systems that had been constructed in the lab. One was a visual programming system oriented around the use of Icons to represent real-world artifacts. Combining and overlapping the Icons were then interpreted as a program that depended on the semantics of the objects; the user was

given various choices to select for right way to do the interpretation. The other was a system called IconicBrowser, which is a browsing front end to a object-oriented database system. Using it, an individual could, for example, retrieve various types of information (including digitized pictures) about all (say) CD players that meet certain defined specs (e.g., price, size). The system runs on two Tandy PCs with the front end on one, and the actual database on another. It was impressively fast I thought, given the hardware limitations, the use of the network, and the transmission of digitized pictures. To a novice like me, both systems seemed very impressive.

While Ichikawa, I and various colleagues were conversing, his poor assistant was trying to find us a place where we could get a quick bite to eat before I left on the train to Hiroshima station. The major problem was that it had to be relatively early (~4.30pm) to get me to the airport in time to catch my plane back to Tokyo. After a few attempts, she found us a place in Seijo, close by. Ichikawa was extremely apologetic about the place before we went, claiming it was primarily a drinking place. I had visions of some dark, dingy, smoke-filled bar, but it turned out to be a very nice and quaint restaurant with small, individual rooms containing tables and chairs. We were joined by Ito-san, a technical staff member in his lab who had received his education at a vocational type school. (As far as I can recall, he's the first such person I've seen.) While we ate, we continued our previous conversations.

Following dinner, Ichikawa dropped me off at the Seijo train station, where Ito-san helped me get on the right train. From Hiroshima station, I took a taxi to Hiroshima airport, where I arrived in plenty of time to catch my plane.

Addresses:

Prof. Tadao Ichikawa  
Dept. of Information Systems  
Faculty of Engineering, Hiroshima University  
Shitami, Saijo-cho  
Higashi-Hiroshima 724, Japan

### **Toshiba Software Laboratory (June 4, 1990)**

Monday was the day I was scheduled to visit the Toshiba Software Laboratory, which is located a 10-minute walk from the JR Kawasaki station. I arrived a little before my appointed time of 1130, and was met at the gate by Ms. Akamine, a secretary for my host Dr. Honiden. The lab is in a new building located on the grounds of a big, post-war Toshiba factory, which, I later found out, manufactures primarily optical scanners for ATMs and the postal service. Upon arriving, I was immediately met by Dr. Honiden and escorted for a meeting room, where I met by Yasukuni Okataku, who is involved with a project for building fault-tolerant distributed systems, and Takeshi Kohno, who is the senior manager. They first gave me an overview of the lab. Basically, it was established a couple of years ago at a very high level in the company hierarchy. The space it occupies is on the 5th floor of the building, and is high quality space organized around an open plan with cubicles in a manner similar to US companies I have visited. It employs approximately 160 people, of which 120 or so are researchers, about evenly divided between BS and MS levels. Their research charter is broad, encompassing the three general areas of systems engineering, AI, and Software Engineering. Their research horizon is supposed to be >5 years, as opposed to the 3-5 yr. horizons for other Labs in the company. Their brochure outlines a number of projects, of which they had chosen two--the IDPS fault-tolerant DOS, and the Mendel Zone concurrent program synthesis system--for later description and demonstration.

Following lunch, I was given lectures and demos on the two systems mentioned above. Both were very relevant to my own areas of interests, perhaps more so than any other projects I had seen in Japan. IDPS is a project to build fault-tolerant distributed systems for applications such as railroad signaling, factory automation, etc. There are three layers: the bottom layer is the DOS, which supports a "fail-stop" broadcast function and "fail-stop" objects, the next contains a distributed file system, database and production system, and the top layer is the application written in an object-oriented fashion. The basic programming paradigm involves replicated objects that interact using the reliable broadcast. Among the unique features they claim are an interesting first-CN-first-serve method for dealing with replicated broadcasts (versus a majority method); among the advantages of this are the fact that you needn't know how many copies there are, which makes it easier to dynamically add and delete members. (Indeed, when I asked them whether they use a membership protocol to agree on failures, they seemed to indicate that they don't need it. Whether this is a by-product of the application or this particular scheme of replication, I don't know.)

The second lecture was a presentation by Naoshi Uchihira on a system that they call Mendels Zone. The purpose of the system is to synthesize concurrent programs from a combination of temporal logic and Petri net specifications. In particular, the temporal logic is used to specify the synchronizing part of the processes, while the Petri nets are used to specify the remainder. These specs are translated by the system into Buchi automaton and a reachability graph, respectively, which are then essentially "intersected" to get the result. The use of the two specification methods is in keeping with the goals of improving both correctness and software productivity; the improvement in correctness is achieved by the synthesis of the difficult part of a concurrent program, the synchronization part, while the improvement in productivity results from the reuse of software facilitated by the Petri net formalism. This work is supported financially by ICOT. I was given papers later on both of the lecture topics.

Demonstrations of the systems followed the lectures. They have set off the ICDP demo as a fairly fancy railroad control system, with an LED panel representing trains and everything. Each object was replicated on each of three processors. To demonstrate its adaptability, they dynamically added a new station to the system without shutting anything else down. They also illustrated the features for fault-tolerance and recoverability by shutting down one of the processors and then restarting it without affecting the running application. I found it all most impressive.

Across the room was the ICOT area, where I was given the demonstration of the Mendels Zone system by Uchihara. It consisted basically of the development of a solution to the Dining Philosophers problem. I again found the demo impressive and convincing, with a fairly sophisticated graphical interface that they said was actually easy to construct using the standard window package. I asked about their view of the feasibility of using temporal logic in real life, and they indicated that their experience had been positive so far. However, Uchihara did say that he thought that verification rather than synthesis was probably a more practical approach to the problem of developing correct software.

Addresses:



Dr. Shinichi Honiden  
Research Scientist, AI Technology Group  
Systems and Software Engineering Laboratory  
Toshiba Corporation  
70 Yanagi-cho, Saiwai-ku  
Kawasaki-shi 210, Japan

Yasukuni Okataku  
Senior Research Scientist  
Systems and Software Engineering Laboratory  
Toshiba Corporation  
70 Yanagi-cho, Saiwai-ku  
Kawasaki-shi 210, Japan

Naoshi Uchihira  
Research Scientist, AI Technology Group  
Systems and Software Engineering Laboratory  
Toshiba Corporation  
70 Yanagi-cho, Saiwai-ku  
Kawasaki-shi 210, Japan

#### **NEC C & C Laboratory (June 6, 1990)**

On Wednesday, I visited the NEC C&C Software Research and Development Group in the Tamachi area of Tokyo. My host was Dr. Toshio Miyachi, manager of the "Basic Systems" department and a graduate of Tokyo Institute of Technology. Upon arriving, we went straight to the 14th floor of the building, where the guest conference rooms were located. There, Miyachi gave us an overview of the structure of the C&C laboratory and the work going on. There are about 300 total people in the lab organized into 3 laboratories, a microcomputer group, and a software planning office. In keeping with NEC's status as the leading producer of Japanese PCs, much if not all of the work is microcomputer-based.

In the basic software lab, there are about 60-70 researchers, half with MS degrees and half with some sort of bachelors' degree (some from 2-year schools.) This particular lab sponsors work in three general areas: basic software for the V and Vr series 32-bit microprocessors, real-time software technology, and advanced software development environments. In addition, the lab is the home of the Unix Technology Center, which acts as a Unix resource for internal NEC purposes, participates in Unix standardization, etc. Most of the work in distributed and fault-tolerant computing is rooted in real-time applications, historically switching applications and more recently space applications (e.g., Japan's space station effort). However, for the latter, they were lamenting the fact that (a) the money devoted to Japan's space effort is small and shrinking, and (b) there is resistance to using 32-bit designs in space because of increased latch-up problems due to radiation. The fault-tolerance work is done as part of the real-time RX-UX OS, which is in turn based on ITRON.

Following these descriptions, we went to the 5th floor, where the laboratory itself is located. There, we were given a demo of the Mustard real-time OS for a prototype multiprocessor based on the V-series chips. The demo was given Shuichi Hiroya, who heads up a team of ~10 people working on the project. The first part of the demo involved running an application on the system

with a graphic display illustrating the status of processes as executing, in the kernel, waiting on a spin-lock to enter the kernel, etc; the spin-lock facility is an example of one of the things that they added to Unix in the process of constructing RX-UX. The second part illustrated that the system could continue executing even though one of the processors was turned off.

The remainder of the visit consisted of my own research presentation, and a nice visit with Yasuyo Kokubo, the general manager of the common software development laboratory.

Addresses:

Dr. Toshio Miyachi  
Engineering Manager, Basic Systems Department  
C&C Common Software Laboratory  
NEC Corporation  
11-5 Shibaura 2-Chome, Minato-ku  
Tokyo 108, Japan

Mr. Shuichi Hiroya  
Basic Systems Department, C&C Common Software Laboratory  
NEC Corporation  
11-5 Shibaura 2-Chome, Minato-ku  
Tokyo 108, Japan

Mr. Yasuyo Kokubo  
General Manager, C&C Common Software Laboratory  
NEC Corporation  
11-5 Shibaura 2-Chome, Minato-ku  
Tokyo 108, Japan

### **University of Tokyo (June 15, 1990)**

On Friday, I visited Prof. Aki Yonezawa's lab in the Department of Information Science at the University of Tokyo. Using instructions provided by his research associate Satoshi Matsuoka, I first went to Shinjuku, then took a Chou line express train towards Tokyo Station, and finally a bus from Ochanomizu station to building number 7 on the university's Hongo campus.

I arrived about 10 minutes prior to my appointed meeting time of 10AM, and was met by Yonezawa, who received his graduate degree at MIT. He had only recently moved to U. of Tokyo, having spent a number of years in the Information Science department at Tokyo Institute of Technology. As a result, he was still somewhat in the process of coming up to speed on establishing his lab.

We chatted for a few minute about his research, which is in the area of object-oriented programming. Specifically, much of his work as been concentrated around the development of an concurrent OO language called ABCL/1. The project has been going on for a number of years, and a book of collected papers was recently published by MIT Press. He also was co-editor with Mario Tokoro from Keio University of a previous MIT Press book on concurrent OO languages. He is also heavily involved with the OOPSLA conference, serving as the PC chairman for the current conference to be held in October in Ottawa.

Following my research talk, we adjourned to lunch. For this, Yonezawa, Matsuoka, and I were joined by Kazuhiko Kato, a research associate from the laboratory of Prof. Masuda who is

interested in databases and distributed operating systems. The conversation centered around the the situation at the U. of Tokyo and the general state of CS in Japan. Among the points raised:

- There were approx. 3 PhD students, 7 MS students, and 6 undergraduates currently associated with Yonezawa's laboratory. They expect the number of undergraduates to rise as they are forced to accept more students because of overall rising enrollments.
- The entire department is small, with only 9 professors.
- PhD students are expected (almost required) to have some journal papers in print before they graduate. As might be expected, this can delay graduation for quite some time.
- They definitely feel the lack of PhD graduates in CS in Japan is hurting them in attracting new faculty. Not being able to support students financially also stymies attempts to expand PhD enrollment.

Following lunch, I said goodbye to Yonezawa, who had an unavoidable meeting. Hence, for the rest of my visit, I was put in the hands of Matsuoka, who shuttled me around showing me various things and taking me to meet with a number of people. Our first stop was Yonezawa's research lab, which contains an impressive array of equipment, including Sun-3s, Sun-4s, Sony NEWSstations, one of the new NEWS MIPS stations, some Symbolics workstations, one of the IBM-Japan's experimental TOPS-1 shared-memory multiprocessors, and an Intel iPSC-2 (the first I'd seen in Japan.) In response of my comments, Matsuoka confirmed that Yonezawa was indeed very successful in obtaining research equipment. In fact, the major limitations have on acquiring equipment have been (1) inadequate A/C (even though the building is relatively new) and (2) space. Indeed, they sometimes have to be selective on workstation usage when the A/C fails.

Next was the departmental instructional lab, aimed at supporting exercises (ie, labs) for undergraduate students from the junior level on down (seniors are associated with labs, and hence, tend to use the lab's research equipment.) The room I was shown contained approx. 8 Sun workstations and 33 X-terminals from a company called JCC (Japan Computer Corp.). The whole setup was relatively new (several months old) and was being rented lock, stock, and barrel from Sun. We chatted for awhile on the virtues of X-terminals vs. workstations, with Matsuoka indicating that terminals were better because they reduced the maintenance requirements, but that they caused a lot of ethernet traffic and required adequate memory on the servers (at least 12-16 MB).

We next went to talk with Kato about his work, specifically the newly established XERO distributed operating system project. Among the interesting features they claim is the use of types and a persistent shared object space, apparently aimed primarily at AI-type applications. To implement the shared aspect, they are developing a lazy caching mechanism. Another feature of XERO is the use of a two-level scheduling mechanism, with tasks being scheduled by the kernel, but threads by the task supervisor within the task. Their claim with respect to this latter feature is that it makes such things as Unix pipelines more efficient than the single-level mechanism in, say, Mach since the shell and the commands can all be threads sharing a single address space. We also talked briefly about the database language work that is apparently intended to be the primary application; he gave me a paper from the most recent ICSDS-10 that describes that research.

In the middle of these discussions, we took a break to get some graphics demos from Toshihisa Shinagawa, a research associate in Prof. Kunii's lab. The demos were mostly on

videotape, and included such things as simulation of trees growing (shown at the Siggraph conference a while back), and people involved in martial arts. There was also one live demo by a PhD student on simulating the way ink spreads by capillary action on a particular kind of paper (no motion on this one yet); although this is outside my area, this is apparently an interesting variant of a canonical problem.

Finally, I spent some time with Matsuoka talking more about research underway in the Yonezawa lab and other things in general. Some of it was about some rather theoretical work on how having synchronization constraints interferes with true inheritance mechanisms. He was especially interested in finding out why my colleague Greg Andrews' SR programming language doesn't have this problem, and after awhile we agreed that it was because (1) in SR, inheritance requires that you replace all of the operation implementations (methods in true OO terminology), and (2) synchronization constraints are not included in the specification. We also talked some about language mechanisms for fault-tolerant programming, and how he thought it would be easy to incorporate the type of event-handling mechanism that I described into my talk into ABCL/1. He also expressed great interest when I mentioned that I was thinking of writing a survey paper on such language mechanisms, and I promised to send him a copy when done.

Addresses:

Prof. Akinori Yonezawa  
Department of Information Science  
Faculty of Science, Tokyo University  
Hongo 7-3-1, Bunkyo-ku  
Tokyo 113, Japan

Satoshi Matsuoka  
Department of Information Science  
Faculty of Science, Tokyo University  
Hongo 7-3-1, Bunkyo-ku  
Tokyo 113, Japan

Kazuhiko Kato  
Department of Information Science  
Faculty of Science, Tokyo University  
Hongo 7-3-1, Bunkyo-ku  
Tokyo 113, Japan

Toshihisa Shinagawa  
Department of Information Science  
Faculty of Science, Tokyo University  
Hongo 7-3-1, Bunkyo-ku  
Tokyo 113, Japan

### **Electrotechnical Laboratory and Tsukuba University (June 28, 1990)**

On Thursday was my day-trip to Tsukuba "Science City" to visit Tsukuba University and the Electrotechnical Laboratory. The former is a relatively new university strong in science, while the latter is the premier national laboratory for research in computer science. When I arrived at Tsukuba Center on the bus from Tokyo Station, I was met by Dr. Yutaka Ishikawa,

who works in the Programming Language section of ETL. He graduated from Prof. Mario Tokoro's laboratory at Keio University about 4 years ago and has worked at ETL since then, although he spent one of those years at Carnegie-Mellon University in Pittsburgh working with Dr. Hideyuki Tokuda on the development of the ARTS real-time operating system kernel.

Ishikawa took me by car to ETL, which are a series of very modern buildings in the AIST research park in Tsukuba. There, he introduced me to Dr. Kokichi Futatsugi, my official host and the head of the Programming Language section. As it was close to the time set for my talk, he gave me only a very brief overview of ETL and his section at this time. The section which is in the Computer Science Division, has 9 researchers, of which 4 currently have a PhD (one more is currently working on it.) He indicated that both the size of the section and the educational level were about average for ETL. As is also normal for ETL, there are no technical support staff at all associated with the section. He said this is one of their biggest problems since it naturally means that the researchers have to do literally everything. One technique they sometimes use to get some technical support is to have long-term visitors from industry. They're there essentially to learn, but are also utilized to do programming, etc. About 1 or 2 at a time seems normal. He also said that students are sometimes used in a similar role.

Following my talk, Futatsugi, Ishikawa, Michiharu Tsukamoto of the Distributed System Section, and I went to lunch at a restaurant in one of the ETL buildings. We spent the time talking about CS in Japan vs. the US, and other general topics. For example, we talked about recruiting people for ETL. Since ETL is viewed as a prestigious place to work in Japan, they did not seem to be having problems hiring PhDs despite a nationwide shortage that I heard about almost every place I visited. The pool they they draw from is the same as the universities -- those inclined basically to an academic-type environment rather than industry -- a fact that I think would be similar to national laboratories in the U.S. In this context, we also talked about the way in which hiring is done at ETL and how it compares with the situation in the U.S. For example, people with an MS degree and wanting a job at ETL have to take a rigorous and competitive exam. I judged this to be somewhat akin to the civil service exam in the U.S., but more more technically oriented and difficult. They are, however, able to bypass this step for candidates with doctorate degrees.

Following lunch, I sat down with Futatsugi for a short one on one talk. He first told me about some of the work going on in his section. His general interests personally are in specification, especially protocol specification similar to the work going on at Tohoku University in Sendai. In particular, one thing he's been working on is a LOTOS interpreter written in OBJ, an algebraic specification language that was done while he was visiting SRI in California. He gave me several papers on the OBJ work. Other projects in the section include a software environment designed to work with algebraic specifications and tailored specifically to the evolution of software, and the work on real-time operating systems and languages being done by Ishikawa.

We also chatted about more general things, such as his experience at SRI and the support programs for foreign visitors at the labs. He also described a new tentative MITI program on New Software Structuring in which he is involved. He made it sound like a very modest program, although he did make a point in saying that it was the first MITI program to concentrate solely on software.

My next visit was with Tsukamoto down in the lab of the distributed systems section. The focus of their work has been primarily on OZ, a distributed object-oriented operating system

based on OSI standards. The OSI aspect is seen to be significant, as is the emphasis on object migration. The work is funded as part of the MITI large-scale project on database interoperability (someone made a comment that anything involved in "information systems" in a general sense fell under the database moniker as far as MITI was concerned.) He gave me a fairly long description of the project, plus a paper. Some of the researchers then gave me a demonstration involving a couple of robot arms moving disks on a platform to the tune of a Towers of Hanoi algorithm.

The next item on my agenda was a visit with Toshio Shimada, who is involved with the SIGMA and EM projects in the Computer Architecture Section. We were running out of time, so this consisted primarily of a presentation on his part about the project. Technically, what I heard corroborated what David Kahaner wrote in his trip report, with the slight exception that Shimada stated that EM-4 was intended for general purpose computing rather than only symbolic computing as implied in DK's report. When I asked about when the various projects started, he said that SIGMA started in 1982 and the EM project in 1986. An interesting side note: the presentation was also attended by 3 US graduate students who were visiting ETL, one from Colorado, one from Stanford, and the third from New Mexico State.

At this point, I left ETL and was driven by Ishikawa to Tsukuba University for my visit there. The campus itself was huge -- I was told several times that it was the largest in Japan (I believe it too.) The building that was our destination was back towards the back of the campus, so we had an opportunity to go through the whole thing. Once there, I met with Dr. Yoshihiko Ebihara, an Associate Professor in the computer science department who had graduated from Prof. Noguchi's laboratory at Tohoku University a number of years ago, and whose current research interest is networks. After introductions, he proceeded to describe the department, or, more accurately, the Institute of Information Sciences and Electronics, giving me a copy of an English brochure as well. It turns out that the department has 51 faculty members, which is huge by Japanese standards and large even by US measurements. (Some people are, however, apparently working in areas that we would consider peripheral to CS.) There are 96 graduate students, with the vast majority apparently being in the MS program; I was told there were about 4/year in the PhD course. There are also about 120 undergraduates/yr in the program, a 50% increase in the normal enrollment of 80/yr due to the recent increased enrollments (noted consistently at other universities as well.)

Following my technical presentation, Ebihara and I continued our previous discussion together with Prof. Ikuo Nakata, the chairman of the Institute. We talked some about the structure of the university, which was established in 1973 more or less on the American university model rather than the koza (chair) model as is common at other Japanese universities. Nakata professed to like the system better, mainly because of the extra flexibility it provides (e.g., came hire into any research area rather than being constrained by the requirements of the chair.) Among the negative points about the system was that it required more meetings among the faculty to resolve issues. He said that the full professors met approximately 2-3 times/month, while the entire faculty of 51 met about once every 3 months. We also talked some about Monbusho support for the Institute, which appeared to be very good. For example, the size of the faculty is large, yet there are still getting additional faculty positions to cope with increased enrollment. We continued this discussion until about 5:30 PM, at which time Ebihara took me to the train station to catch the next train to Tokyo.

Addresses:

Prof. Ikuo Nakata  
Institute of Information Sciences and Electronics  
University of Tsukuba  
Tsukuba, Ibaraki 305, Japan

Prof. Yoshihiko Ebihara  
Institute of Information Sciences and Electronics  
University of Tsukuba  
Tsukuba, Ibaraki 305, Japan

Dr. Koukichi Futatsugi  
Chief, Computer Language Section  
Computer Science Division  
Electro-Technical Laboratory, AIST-MITI  
1-1-4 Umezono, Tsukuba  
Ibaraki 305, Japan

Dr. Yutaka Ishikawa  
Computer Language Section  
Computer Science Division  
Electro-Technical Laboratory, AIST-MITI  
1-1-4 Umezono, Tsukuba  
Ibaraki 305, Japan

Michiharu Tsukamoto  
Senior Researcher, Distributed Systems Section  
Computer Science Division  
Electro-Technical Laboratory, AIST-MITI  
1-1-4 Umezono, Tsukuba  
Ibaraki 305, Japan

Toshio Shimada  
Chief Scientist, Computer Architecture Section  
Computer Science Division  
Electro-Technical Laboratory, AIST-MITI  
1-1-4 Umezono, Tsukuba  
Ibaraki 305, Japan

**Shinshu University (July 2, 1990)**

On July 2, I visited the Faculty of Engineering campus of Shinshu University in Nagano City. My host for the visit was Prof. Masayuki Okamoto of the Department of Information Engineering. Prof. Okamoto is a former student of my host scientist at Tokyo Institute of Technology, Prof. Takuya Katayama, and had been at the department at Shinshu since its founding in 1978 or so. After arriving at the university, he first gave me a brief rundown of the department. There are currently 7 chairs in the department, which is organized like most Japanese universities into the chair system (koza in Japanese.) In his lab, there is currently one

associate professor (working in theory of computing), and about 15 MS students. There is currently no PhD program in the department. He said that they have proposed one, but that the Ministry of Education (Monbusho) is balking at the cost despite the current shortage of Ph.D.'s in computer science.

After this brief introduction, we adjourned to a lecture room on another floor for my talk. He had indicated when we were setting up the visit that a research talk was probably too specialized, so I had volunteered to give an "Introduction to Distributed Systems" talk based on a lecture I had given earlier in Tokyo. The turnout was impressive -- about 50-60 people -- of whom the vast majority were students.

Following the talk, he spent some more time showing me around the department as we talked more about his research, the department, and the university in general. His basic work is involved with pattern recognition algorithms, specifically, the construction of document recognition systems. He had his students show me a demo of the latest system, which will convert pages from a scanned scientific document into TeX input. The system, which is written in C for Sun 3/4 workstations and has an X-windows interface, deals with both roman and Japanese characters, tables, scientific equations, and figures. I certainly not an expert in such systems, but it compared quite favorably to the commercial systems that I had seen. He gave me a copy of a paper that he had just returned from presenting at an IEEE Document Recognition Workshop in New Jersey. He mentioned also that people at Bell Labs are working on such systems and are perhaps the leaders in the field. By the way, the laboratory itself contained a fairly impressive array of equipment, including Suns, NEC PCs, and an Omcron Luna workstation.

Following the demo, he showed me some of the departmental instructional machines. The "big iron" here was a Convex that runs Unix and can support up to 200 undergraduate users simultaneously. The machine was bought 3 years ago using money provided by Monbusho, who also supplies money for maintenance. In response to my query, he says that Monbusho will, as a rule, give money for such equipment about once every 10 years. Since this time period greatly exceeds the life span of most computer equipment, they have to be careful about being stuck with obsolescent equipment. The strategy they use at Shinshu is to lease the equipment so that it is more easily replaced. When I asked about why they had chosen a Convex, he indicated that the availability of NFS at the time of purchase was one important consideration.

We then went up on the roof to admire the view (the building is one of the tallest buildings in Nagano.) He also showed me the tennis courts next door which are scheduled to be replaced with a new building within the next few years. The building be about 50% of the size of their existing building, and is being constructed mostly to help accommodate their increased enrollment. Since Okamoto enjoys tennis, I got the feeling that he had decidedly mixed emotions about the whole thing!

Upon returning to his office, he gave me a brochure describing the department. Although it is in Japanese, it does contain an understandable picture of the departmental network. There is also a university network developed largely by the department. They are connected to JUNET through a line to an Epson facility in Matsumoto, where the main campus of the university is located.

I came away impressed with the situation at Shinshu. For a university at the top of the 3rd tier in Japan (as Okamoto characterized it), the equipment and general Monbusho support seem



better than I had expected. I was also impressed that Okamoto himself maintains an active research program and is successful in publishing. I frankly can't think of many people at equivalent universities in the U.S. that do as well. While his kind of success may not in fact be the norm, my overall impression is that there is perhaps less distance between the best universities and universities such as Shinshu than one would find in the U.S.

Address:

Prof. Masayuki Okamoto  
Dept. of Information Engineering  
Faculty of Engineering, Shinshu University  
500 Wakasato, Nagano 380  
Japan

### **Hitachi Systems Development Laboratory (July 4, 1990):**

On Wednesday, I visited Hitachi at the invitation of Dr. Hiro Ihara, a person I've known for a number of years by virtue of our common interest in fault-tolerant computing. Although Ihara currently works for Hitachi's relatively new Space Systems Divisions, he arranged for me to visit Hitachi's Systems Development Laboratory since the work there is generally more in tune with my own interests. He has also indicated that he'd like to show me some of the other labs, an offer that I intend to take him up on during some return visit.

I arrived at the Tama Plaza station on the De'en Toshi line on the west side of Tokyo about 2:00pm, where I was met and taken to the labs by taxi. Upon arrival, I was led to a conference room, where I sat for few minutes before being greeted by Dr. Kinji Mori, a senior researcher who is currently director of the ADS (Autonomous Decentralized System) project. We talked briefly about the lab and his work, and then saw a short slide show giving an overview of the work going on at the lab. There were a number of interesting projects underway, with the two most relevant to my own work being ADS, a distributed and decentralized system for real-time process control, and the fuzzy logic based subway control system.

After making my research presentation, we returned to the conference room, where I was shown a videotape on ADS. It was interesting and informative, despite being mostly non-technical. The tape progressed by making the analogy between the cooperating processors in ADS and the cells in the human body. I originally thought this was mostly hype, but interestingly enough, it turned out to be historically accurate -- the work on ADS was originally motivated by conversations with an MD turned computer advocate at the University of Tokyo who made this analogy!

The videotape and subsequent technical discussions highlighted a number of interesting technical aspects of ADS, especially in the use of decentralized algorithms and what they call a "Data Field" for communication between processes. Although I have not yet sorted out all of the technical details, the DF concept seems very reminiscent of the tuple space found in the Linda system developed by David Gelertner, et al at Yale. Although the DF was probably developed prior to Linda (the ADS project started in 1977, with the first patent in 1979), it is not nearly as well publicized among academics (as Mori remarked somewhat ruefully, he actually spends most of his time writing patent applications rather than papers.) He did give me a paper on the software structure of ADS that appeared in 1986 Fall Joint Computer Conference in Dallas, which has pointers to a number of other earlier papers as well. (He also said that there is a more technical

videotape on ADS, although it is unfortunately only in Japanese.) From a commercial standpoint the system has apparently been very successful in the real-time process control arena for which it was originally designed. They have also been successful in obtaining a number of patents worldwide on this decentralized control scheme (including the US), receiving some sort of yearly award for best patent in Japan in the process. My impression is that this is a prime example of a very worthwhile and interesting project that is virtually unknown in the US.

Ihara and Mori then accompanied me into another room where a stock trading optimization program based on neural nets was demonstrated. The person doing the demonstration was Dr. Ikuo Matsuba, a senior researcher who I believe is head of this project. The basic thrust of the program was to select 5 stocks out of 100 so as to maximize the return but minimize the risk. The basic strategy seems to be to represent historical correlations (both positive and negative) between pairs of stocks as neural links, and then iterate to determine which neural connections strengthen as time progresses and which connections weaken. The strengthened links indicate "good" combinations, while the weakened connections indicate "bad" combinations. The result could then be used to select the particular combination of stocks that will maximize the return for a given level of risk. To demonstrate that this was a good choice, they plotted this curve on a graph of risk vs. return, and then selected 5000 (?) combinations at random and plotted them; all 5000 points were below the curve suggesting that optimal combinations had indeed been chosen.

The demo was run on a PC level machine, although the calculations had been actually run beforehand on a supercomputer. The research group has also successfully fabricated a chip specifically designed for doing this type of neural net calculations. Given all this, I was surprised to learn that the number of people working on this in the whole corporation was in the range of 20-30. The particular program I was shown seems to be designed solely to demonstrate the neural net approach, and I got the feeling that there were no plans to commercialize it (despite a lot of jokes about getting rich using the program as advisor.)

At this point Ihara, Mori and I adjourned to dinner at a nice combination French/Japanese restaurant. The discussion during dinner ranged from the usual discussions of Japan and the US, to discussions about Hitachi's efforts in space-based computers. According to Ihara, they started 3 years behind other Japanese companies in this area, but are aiming to be competitive in the early 21st century. This work is being done in partnership with TRW in the US.

Addresses:

Dr. Hirokazu Ihara  
Senior Chief Engineer, Space Systems Division  
Hitachi, Ltd.  
216 Totsuka-machi, Totsuka-ku  
Yokohama 244, Japan

Dr. Kinji Mori  
Senior Researcher, 1st Research Dept.  
Systems Development Laboratory  
Hitachi, Ltd.  
1099 Ohzenji, Asao-ku  
Kawasaki-shi 215, Japan

Dr. Ikuo Matsuba  
Senior Researcher, 1st Research Dept.  
Systems Development Laboratory  
Hitachi, Ltd.  
1099 Ohzenji, Asao-ku  
Kawasaki-shi 215, Japan

### **IBM Tokyo Research Laboratory (July 5, 1990)**

The day following my visit to Hitachi, I traveled to IBM TRL for my last visit during this stay in Japan. The arrangements were made with Dr. Norihisa Suzuki, the director of the lab. David Notkin from the University of Washington was also arranging a visit at the same time, so we ended up going together and giving back-to-back talks. Since David had visited the previous year and had research interests more in line with the work going on there, I had the luxury of sitting back and listening more than on some of my other visits.

David and I arrived at the lab at about 12:45, where we were met by Dr. Toru Takeshita. Dr. Takeshita is Manager of the Computer Science Institute of the lab and the author of a recently completed book (in Japanese) on CASE (Computer Assisted Software Engineering.) We chatted in the lobby about various things for about half an hour (including how important Japan is to IBM -- apparently it accounted for 10% of total worldwide profits in the most recent fiscal year) , at which time we headed to a nearby building where the actual talks were to be held. My talk was first, and lasted about an hour; Dr. Suzuki entered a few minutes before it finished. David talked second, and put us back on schedule by ending in about 45 minutes.

The next thing on our agenda was a series of discussions and presentations. The first thing was a presentation by Tamiya Onodera on the COB programming language project. (The head of the project was out of town.) COB is an object-oriented version of C that solves some of the problems associated with C and C++. As was remarked several times, (by both them and us), they are probably fighting a somewhat losing battle despite a better product due to the widespread use of C++; as they said, they really started a couple of years too late to make a big impact. Nevertheless, they have some interesting results and were able to give a nice demonstration of a development environment system for COB including a debugger that shows a dynamic graphical representation of the program during execution. As David remarked, the debugger is a nice example of a different kind of program animation.

We then returned to the conference room, where Takeshita gave us an overview of the TRL and its place within the IBM corporate structure. The number of researchers in the lab is about 260, which compares to 2400 at TJ Watson and 750 at Almaden. The four specific areas covered by TRL are Computer Science, Computer Applications, Component Technologies, and Manufacturing Technologies. He said that TRL is considered especially strong in the applications field. In that area, they currently have efforts in graphics, application development environments, man-machine interfaces, natural language processing, image processing, knowledge-based systems, and prototyping. On natural language processing, he indicated that their English to Japanese translator is unique since it first rearranges the English into a more Japanese word order, and then translates to Japanese. The result, he said, is 70% good and is currently used for translating such things as technical manuals. In the CS area, there are projects in systems (architecture, OS, programming languages and environments, software tools, the TOP-1 multiprocessor, and multiprocessor AIX) and theory (theoretical CS and OR.) Takeshita

then gave us a talk on a description and history of Japanese software development efforts, with an emphasis on CASE.

The next talk was a description of the TOP-1 multiprocessor project. This presentation was made by Dr. Nagatsugu Yamanouchi, who is manager of the workstation software group. The TOP-1 project, which is well-known in Japan, started in 1987, with the hardware prototypes being finished in 1988 and the OS in 1989. The machine is a shared memory multiprocessor consisting of 10 386-type multiprocessors. Standard memory configuration is 32 Mb, but it can take up to 128 Mb. Claimed as one of its unique features is its snoopy cache, which consists of 128 kB of storage per processor; the cache uses a dynamic variant of standard snoop protocols, and is switchable between either of two modes on a per processor basis. Another unique feature is that the hardware is asymmetric since the disk subsystem is only accessible to only one processor. As a result, the OS has also been constructed asymmetrically, with the one processor that can communicate with the disk dedicated to executing the kernel and the others running user tasks. Also, the system has been designed with a built-in hardware monitor to facilitate collection of various statistics, something that has not been done before with this kind of architecture. The machine was originally targeted for numerical applications, although it doesn't really enough horsepower to be useful currently and the asymmetric hardware and software architecture would seem likely cause bottlenecks if the number of processors were expanded significantly. Two languages are currently being used: a Multiprocessor Lisp developed at TRL and a MP Fortran that is apparently based on the IBM parallel Fortran developed at Yorktown Heights. Ten prototypes of the system were built, with four currently on loan to Japanese universities (specifically, Keio, Waseda, Tokyo, and Kyoto.)

At the end of all this, we were taken to Suzuki's office, where we sat and chatted for awhile. He talked quite easily and freely about the projects in his lab, and seemed very interested in the comments that we offered about the COB and TOP-1 projects. He clearly supports the type of basic research represented by these projects, but with the recognition that part of his job is to ensure that the company's money is used productively.

Addresses:

Dr. Norihisa Suzuki  
Director, IBM Tokyo Research Laboratory  
5-19 Sanbancho, Chiyoda-ku  
Tokyo 102, Japan

Dr. Nagatsugu Yamanouchi  
Manager of Workstation Software  
IBM Tokyo Research Laboratory  
5-19 Sanbancho, Chiyoda-ku  
Tokyo 102, Japan

Dr. Toru Takeshita  
Manager of Computer Science Institute  
IBM Tokyo Research Laboratory  
5-19 Sanbancho, Chiyoda-ku  
Tokyo 102, Japan