

REPORT ON A VISIT TO ICOT

10~~X~~ June to 28 June, 1985

Douglas S. Kerr

Department of Computer and Information Science
The Ohio State University
2036 Neil Avenue Mall
Columbus, Ohio 43210
USA

1.0 INTRODUCTION

I was invited to visit ICOT to discuss the problems of evaluating the performance of database machines in general and the particular relational database machine DELTA developed by researchers at ICOT. This mutual interest in database machines began a couple of years ago when Mr. Takeo Kakuta and Mr. Yoshimitsu Hirai of ICOT visited the Multi-Backend Database System project at Ohio State University. The interest was continued when Mr. Haruo Yokota of ICOT visited the Naval Postgraduate School during the 1983-84 academic year while I was visiting on sabbatical leave. I have been following the development of DELTA for a number of years and was very interested in visiting ICOT now that the prototype is working.

Since my interests are mainly in the database area and not in artificial intelligence, my visit also provided an opportunity for me to learn about some of the database requirements for real artificial intelligence applications. It is clear that there are still many problems that must be solved before the large amount of data required by real applications can be accessed efficiently.

2.0 THE RELATIONAL DATABASE MACHINE - DELTA

The first part of my visit consisted in a review of the design of the relational database machine DELTA. This machine has been designed to serve as a backend storage system for the knowledge based systems being developed by the other groups at ICOT. Thus its design is quite different from other database machines.

I was sorry that my visit to ICOT had to be delayed until June since many of the developers of DELTA had returned to their home companies. Nevertheless I did get a good understanding of the design of the machine from Mr. Kakuta, one of the original designers.

2.1 A Demonstration Of DELTA

Although I did not see a "live" demonstration of DELTA, I did see a copy of the demonstration prepared for the ICOT Open House held after the FGCS '84 in November. The demonstration, which involved supplying information for a tourist about places to visit around Tokyo was very interesting. It showed the capabilities of DELTA well.

Before the demonstration I was shown the actual prototype system including the special purpose sorting hardware that is an important part of DELTA.

3.0 THE MULTI-BACKEND DATABASE MACHINE - MBDS

I made a presentation on the Multi-Backend Database System - MBDS, which is under development at Ohio State University and the Naval Postgraduate School. This system is designed to provide enhanced performance using existing hardware configured in a novel way. It is also designed to allow easy expansion. The system has been working for about a year now and we are getting ready to perform extensive performance evaluation experiments.

4.0 PERFORMANCE EVALUATION

The main purpose of my visit was to discuss the problems of performance evaluation of a database system. A small amount of work has been done in this area, mainly at the University of Wisconsin and at the Naval Postgraduate School. However the earlier work differed from what is needed for DELTA in two ways. First, DELTA is intended for use in a different environment, namely a knowledge-based environment. Thus both the type of data stored and the type of processing required will be somewhat different from that assumed in a standard relational database system environment. Second, DELTA stores its data in a novel way, because of the knowledge-based environment, and therefore the previously proposed benchmarks are not really going to be able determine the performance of DELTA under varying conditions. Thus it will be necessary to develop a new set of benchmarks which are more applicable to DELTA and its environment.

5.0 KNOWLEDGE BASED MACHINE GROUP - INTERFACE TO DELTA

Possibly the most interesting part of my visit was learning about the efforts to interface inference systems based on Prolog with a relational system, DELTA. I had several interesting discussions with Mr. Yokota on this subject. I think this interface is one of the biggest problems that must be solved before real knowledge-based applications can be built. Such applications will require the capability of accessing a relatively small amount of data from the large amount of data available. Most applications built so far have

been "toy" applications in the sense that they use only a very small amount of data. In such an environment there is no "database problem", since almost any search method can be applied to find the data that is relevant to a particular processing problem. However when the amount of data available is large, a gigabyte or more say, then there is a "database problem". In this case it will be useless to have a very fast inference machine if it cannot be supplied data at a fast enough rate. It will not be easy at all to locate the small amount of relevant data amongst the large amount of irrelevant data. Thus it is very important that this problem of data access be addressed at the same time as the more theoretically interesting one of "knowledge processing".

One of the discoveries already has been that the standard operations provided by a relational database system are not sufficient for a knowledge-based machine since they require too large a transfer of data between the backend database system and the front-end inference system. Thus Mr. Yokota has been leading an effort to determine what new operations can be added to the database system to allow that system to further reduce the amount of data that has to be transferred to the inference processor. The development of a deductive database system based on unit resolution seems to be a step in the right direction. It has been well established in the database area that a request to a database machine must be to process a large amount of data if the use of a separate database machine is to be effective. Previous efforts to interface with prolog required too much interaction between the inference machine and the database machine. The analogy in the data processing environment would be a host-backend system that required the host to issue a distinct request for each record that it needed to process rather than to issue a single request that would allow the backend to find and transmit a large group of records at once. The former approach is well-known to be unsatisfactory because the limited bandwidth between the backend and host can never be sufficient to avoid a serious bottleneck at this connection.

6.0 KAISER GROUP AND A KNOWLEDGE ACQUISITION SYSTEM

Mr. Susumu Kunifuji of the Kaiser group and Mr. Taizo Miyachi, who recently returned to Mitsubishi from the Kaiser Group gave me an interesting overview of their knowledge acquisition system. They described the overall design of Kaiser and discussed the demonstration that was developed for the ICOT Open House in November. It will be interesting to see how this system works when there is a large knowledge-base already in existence.

7.0 A VISIT TO NTT

Mr. Masanobu Higashida arranged a visit for me to the Yokosuka Electrical Communication Laboratory of NTT. Mr. Higashida, who is now at the corporate headquarters of NTT, spent a year visiting the

Multi-Backend Database System project at Ohio State a couple of years ago and it was good to see him again. During this visit Dr. Akihiro Hashimoto, the Director of the Data Processing Development Division, gave me an overview of the activities at the Yokosuka Laboratory. I then gave a presentation on the Multi-Backend Database System. This presentation prompted several interesting questions about our design and how the system might perform.

After a delicious lunch, I visited the Data Base Management System Section and heard a presentation on the Advanced Database Machine (ADAM) for which a prototype is being developed. This prototype consists of three major components, an intelligent disk controller, which includes a special purpose search processor, a general purpose CPU for control, and a database operation accelerating processor (DAP), which contains a special purpose filter and sort-merger for efficient join processing. The discussion which followed centered on the actual data storage organization and the use of the machine in an information retrieval mode alone. After this discussion I was shown a demonstration of the intelligent disk controller which is now operational. It will be interesting to see how this machine performs when the other components are completed.

My final visit at the Laboratory was with Mr. Toru Ishida of the Knowledge Base Systems Section and concerned their efforts to develop both basic tools for knowledge base management, a knowledge base management system itself and various application systems. Our discussion centered on the need to have a close link between the knowledge processing system and the database system if there was not going to be a major bottleneck between these two systems.

8.0 INTERACTION WITH YOSHIMITSU HIRAI OF JAPAN INFORMATION PROCESSING DEVELOPMENT CENTER

I first met Mr. Hirai several years ago when he and Mr. Kakuta visited Ohio State. Although Mr. Hirai is no longer at ICOT, he visited several times during my stay and we had several interesting conversations. First we discussed the need for the development of an educational program in the use of microprocessors. They require a careful integration of both hardware and software concepts. Such programs do not really exist at this time especially in Japan. They are only recently coming into existence in the United States. Second, we discussed the feasibility of network, particularly mail, connections between Japan and the United States. I hope I provided him with some useful information particularly on the activity of the IEEE Computer Society in the area of computer science and engineering education, which is most relevant to the problem of education in the use of microcomputers.

9.0 CONCLUSIONS

I learned a great deal during my visit, both about the DELTA

database machine and about the efforts of the Fifth Generation Computer Systems Project in general. I think the idea of having a large number of researchers all focusing on one particular large problem is very interesting. This situation forces the solution of all the problems that are interrelated, not just those that seem most interesting or tractable at the moment. I think one of the important offshoots of this approach will be that at least some solutions to the "database" problem in artificial intelligence will be solved. As a "database person", I think this problem must be solved before any real applications can be developed. And, as the Knowledge Base Machine group has learned, the problem is not an easy one to solve. I am certainly glad it is being worked on in conjunction with other problems. I do not think it is easily solvable after all the other problems are solved.

As several other visitors to ICOT have mentioned, I think future collaboration between Japan and the United States would be enhanced by the connection of ICOT to one of the networks used by the US computer science community, either CSnet or even the Unix network called Usenet. There actually seems to be interest in such network connections at other locations in Japan. For example, Mr. Ishida at the Yokosuka Laboratory of NTT also seemed interested in network connections to the United States.

10.0 ACKNOWLEDGEMENTS

My visit to ICOT, which was also my first visit to Japan, has been most pleasant. I would first like to thank Dr. Kazuhiro Fuchi for the invitation to visit and for the generous financial support that allowed me to accept his invitation. The environment he has established at ICOT is an excellent one. I would also like to thank Dr. Hidenori Itoh, Chief of the newly formed Third Laboratory who served as my gracious host. He was most considerate during my visit.

In addition I would like to thank Mr. Higashida for arranging my visit to the Yokosuka Laboratory of NTT and making sure I got there without difficulty.

Many people were most gracious to me during my visit, helping a foreigner who knows no Japanese, but who likes all kinds of food to find delicious food to eat and generally to have a good time outside working hours. In particular, Dr. Itoh and the members of the Third Research Laboratory treated me to my first Japanese dinner and party, which included my first sampling of shochu as well as some very good food. Dr. Fuchi also treated me to an excellent banquet. I still remember the delicious crab which was served. Mr. Hirai took me on my first trip away from the ICOT neighborhood, taking me to Asakusa for another delicious dinner. It was there that I learned sake is served in a small wooden box as well as the small cup usually seen. Mr. Higashida treated me to my first sushi restaurant. Watching the sushi being prepared was a particularly interesting experience. Mr. Shigeki Shibayama and his wife graciously hosted me to a very nice dinner at their home during my first weekend alone in Tokyo. I

greatly appreciated the opportunity to visit in a Japanese home. Finally I would especially like to thank Mr. Yokota, who took special interest in my needs. His hospitality was so great it cannot be enumerated, but I would like to mention at least a few ways he was more than helpful. He saw to it that I got lunch during my time at ICOT. He came to Shibakoen station to escort me to Mr. Shibayama's home for dinner. This trip, which was far out of his way, must have taken a large part of his Sunday. In addition he showed my wife, daughter and myself the fun of an okonomiyaki restaurant. His assistance made my stay in Tokyo outside ICOT a delightful one.

Finally, I would particularly like to thank Mr. Hiroyuki Kusama, who handled all the details of my visit with an incredible amount devotion. His assistance was most helpful and included an attention to detail that I had never experienced before and will probably never experience again.

Douglas S. Kerr

Short Vita

PERSONNAL DATA

Residence: 2104 Iuka Avenue, Columbus, Ohio 43201
Telephone: (614)299-3887

Office: Department of Computer and Information Science
The Ohio State University, 2036 Neil Avenue Mall
Columbus, Ohio 4321
Telephone: (614)422-1519

Birthplace: Washington, D.C.

Birthdate: November 19, 1940

Family Status: Married with one child

EDUCATION

BA Yale University Mathematics, 1962

MS Purdue University Mathematics, 1964

PhD Purdue University Computer Sciences, 1967

EMPLOYMENT RECORD

June 1971 - Present Associate Professor of Computer and Information Science, The Ohio State University

June 1983 - August 1984 Adjunct Research Professor of Computer Science, Naval Postgraduate School

October 1980 - June 1983 Director, Laboratory for Database Systems Research

October 1967 - June 1971 Assistant Professor of Computer and Information Science, The Ohio State University