

Report on a Visit to ICOT

Larry Henschen

Electrical Engineering and Computer Science
Northwestern University
Evanston, Illinois
60201 USA

1. Introduction

I was invited to visit ICOT for a period of three weeks from Monday, September 20 to Friday, October 10. The main purpose of the visit was to work with researchers in the Third Laboratory on topics relating to deductive knowledge bases. Specifically, I was to present some of my own recent results on indefinite inference in deductive databases and was to work with Drs. Itoh, Murakami, Morita, Miyazaki and Yokota on topics relating to the design of the knowledge base systems. A second purpose for visiting ICOT was to become familiar with as many of the research areas as possible. Such visits as mine are, I believe, an invaluable way for ICOT to provide direct and detailed information to relevant scientists around the world; there is a much higher transfer of information and ideas than is possible, for example, at conferences and workshops. ICOT is to be applauded for its commitment to the dissemination of information.

2. Presentations at ICOT

On Friday, September 26, I gave a lecture to the Working Group on Knowledge Base Management on my work on indefinite inference in Non-Horn Deductive Databases. Although this topic is not one of the main research topics presently studied in the Third Laboratory, the audience seemed very interested and attentive and asked some interesting questions, both specific to my research and general to the work of the Third Laboratory.

Most of the first week was taken up with informal presentations by researchers in the different laboratories. These informal discussions are an excellent way for a visitor to get an overall view of ICOT's goals and projects. I found them very informative, and hope that I was also able to make some useful comments to a few of the researchers even though their projects were not directly part of the Third Laboratory.

Toward the end of the first week I also received some presentations by my colleagues in the Third Laboratory. These, of course, were the most interesting to me. The presentation on Friday of the first week by Drs. Murakami and Morita spawned some ideas that led to research during the second week that I hope will be very profitable for ICOT as well as for myself. I continued to receive detailed presentations on the work of the Third Laboratory during the second week, all of which I continued to find most interesting.

My only suggestion concerning these presentations is that it might be more useful to have the detailed presentations of topics directly connected to the visitor and his laboratory first and have the overview presentations later in the visit. This would allow a little more time for collaborative research. In my case, I already knew the work of the Third

Laboratory quite well because of visits to my university by Drs. Yokota and Murakami. A different visitor who did not have this advantage might be somewhat more anxious to "get down to work." On the whole, however, my visit was a great success, as will be seen by my remarks in the next section of this report, and I thoroughly enjoyed my visit both professionally and personally.

3. Collaborative Research - Present and Future

At the end of the first week, Drs. Murakami and Morita discussed with me their work on retrieval-by-unification, specifically their proposals for unification joins and selects and related hardware designs. Although their work is in a very early stage I found it interesting. One aspect of the unification join process caught my attention and sparked an idea. The problem in retrieval by unification is that a large number of unification tests must be performed. Researchers in the Third Laboratory have already recognized this problem and proposed various techniques for dealing with it. In particular, they have shown that an extension of the generality partial order to a total order can be used to significantly reduce the number of unification tests required for a join or select operation. However, I noted that the generality partial order itself could be used in addition to the total extension to replace a (hopefully large) number of unification tests by a simpler and faster computation. This idea is based on the following two remarks. If t_1 is an instance of t_2 , then t_1 and t_2 are unifiable; if t_1 and t_2 are unifiable and s_i is more general than t_i for $i=1,2$, then s_1 and s_2 are unifiable and their MGU can be computed directly from the MGU of t_1 and t_2 . I developed the required algorithms and described some of the storage and hardware consequences of this new approach. However, we have only begun to understand the consequences of this idea as it relates to other aspects of knowledge base management software and hardware. I think this is a problem where collaboration in the near future will be very useful. I hope that my extensive background in theorem proving will allow me to provide some additional insights into techniques for knowledge bases.

On Thursday of the second week, I met with Dr. Yokota, presently of Fujitsu and formerly of the Third Laboratory. His efforts to implement various theorem proving strategies within the knowledge base scheme is extremely interesting to me because my background for this area comes from my previous 15 years work in the area of automated deduction. He described several different theorem proving strategies which he could implement within the retrieval-by-unification paradigm. He was concerned about the refutation completeness of one of these strategies, a particularly important one. At first, I thought this strategy was NOT complete. However, in the process of searching for a counterexample we decided that the strategy might be sufficient after all. I was then able to map out a proof of its completeness, which I wrote in detail that afternoon. I also described for Dr. Yokota a general technique for attacking such completeness proofs that I hope will be useful to him in the future.

I plan to continue collaborating with members of the Third Laboratory. We are implementing at my university a deductive database (a precursor of deductive knowledge base), and many of the problems are closely related. I think it will be useful for both projects to share our ideas and experiences. I hope that the primitive ideas on generality begun during my visit can be fully developed and perhaps reported in the literature.

One suggestion I have for future work at ICOT is to develop and then implement at least one real-life, large-scale knowledge base. Such a

development is crucial to the whole field of knowledge bases in general because we must prove to the general community that our ideas really can help solve present and future real-world problems: we must prove that our systems are not just interesting scientific artifacts. Of course, I am deeply committed to the idea that information systems of the future will be exactly the kind of knowledge bases that I and my colleagues at ICOT are developing. A demonstration system that solves a real need now will be a great impetus for our field. In addition, there is a very practical benefit of having such an example. In many situations in artificial intelligence, it is difficult or impossible to provide a concrete analysis of our techniques. Moreover, often the problems involved with solving AI applications are not even evident until large-scale examples are tried; it is widely accepted now that experimenting with "toy" problems does not lead to any real insight. Having a sample knowledge base of the size anticipated for the future will be an invaluable tool for developing techniques and testing ideas. A knowledge base with more than a thousand rules in addition to a sizable extensional database would be excellent. (Perhaps one could start with a data intensive expert system, that is, one in which there is a large volume of "facts".) The development of such a system is out of the question for a small research setting like a university. However, ICOT with the help of its supporting computer companies does, I believe, have the resources to develop one or more such systems. Such development would be an important step forward for research at this point in our research efforts.

4. Acknowledgements

It has been a great pleasure to work with my friends and colleagues in the Third Laboratory. I have learned a great deal about ICOT in general and about deductive knowledge bases in particular. I hope also that I have been able to make some small contribution to the work of the Third Laboratory and that I can continue to do so in the future.

I wish to thank Dr. Fuchi for approving my visit. I am also grateful to Dr. Itoh, Chief of the Third Laboratory, for his efforts in arranging my visit and for finding so many interesting places to eat lunch. The researchers of the Third Laboratory gave me many interesting presentations which I greatly appreciate. They also gave me and my mother, who was accompanying me in my travels, a wonderful sushi party during the first week: it will remain as one of the most memorable events of my trip.

Mr. Kusama deserves a special vote of thanks for all the arrangements he made (which were excellent) and especially for the concern that he showed for the well being of my mother. I am most grateful to him for his part in making our stay in Japan safe, comfortable and enjoyable.

Finally, my special thanks go to my friends Haruo Yokota and Musaki Murakami for suggesting that I might visit ICOT when they were at my university last February. They have generously shared their problems and ideas with me over the past few years, and I have enjoyed my association with them. I very much hope that it continues.

LAWRENCE J. HENSCHEN

DATE OF BIRTH: October 11, 1944

EDUCATION: B. A. 1966, University of Illinois at Urbana, graduated with honors.
M. A., 1968, University of Illinois at Urbana.
Ph.D., 1971, University of Illinois at Urbana.

EMPLOYMENT: 1985-present, Professor of Electrical Engineering and Computer Science, Northwestern University; Evanston, Illinois.
1977-1985, Associate Professor, Electrical Engineering and Computer Science, Northwestern University, Evanston, Illinois.
Jan.- March, 1973 and June-August, 1974, Argonne National Laboratory, Argonne, Illinois, Visiting Scientist.
1971-1977, Assistant Professor, Computer Science, Northwestern University, Evanston, Illinois.
1966-1971, Research Assistant, Computer Science, University of Illinois at Urbana, Urbana, Illinois.
1964-1966, Grader, Computer Science, University of Illinois at Urbana, Urbana, Illinois.