A report on a visit to ICOT: The Institute for New Generation Computer Technology

Randy Goebel

Logic Programming and Artificial Intelligence Group
Department of Computer Science
University of Waterloo

1 Introduction

From September 2, 1985 to September 27, 1985 I visited the Japanese Institute for New Generation Computer Technology (ICOT) in Minato-ku, Tokyo, Japan. I had been the recepient of a previous invitation in 1984, but I was unable to accept that invitation as I had already planned a scientific expedition to Japan for the only time period that my teaching duties would allow. Negotiations with Dr. Koichi Furukawa resulted in another invitation for 1985, which I gratefully accepted.

There were several reasons for accepting the invitation to do collaborative research. A major reason is my longstanding interest in the use of logic as a foundation for doing artificial intelligence (AI) research. ICOT's choice of logic programming as the basis of its ten year Fifth Generation Computer Systems (FGCS) project was consistent with my personal views of both the basis and methodology of AI research.

In addition, I felt it was both an honour to be invited, and a rare opportunity to find out how research at ICOT was actually conducted. Doing AI research in Canada meant that I was familiar with the "Fifth Generation hysteria" that has been induced in the United States, and I was most eager to find out how the Fifth Generation project was progressing. Furthermore, my previous visit to Japan made me curious about how a new basic research institute, apparently novel for Japan, would be organized and run.

Finally, my close colleague Maarten van Emden had accepted a similar invitation earlier in the year, and had convinced me that the atmosphere at ICOT was conducive to research– one was free from the typical duties of a university professor, whose research time becomes quite fragmented by teaching and administration.

2 Collaborative research at ICOT

My research time at ICOT was to be spend in the Firstf Research Laboratory, whose research was on "basic software systems" for the FGCS project. The chief of this lab is Dr. Koichi Furukawa.

Although it was suggested that I consider possible topics for collaborative research at ICOT before my visit, I arrived with only a few vague ideas. In addition, I had not yet completed a paper that I had been invited to submit for the special knowledge representation issue of New Generation Computing. As Dr. Furukawa was also the editor of that special issue, I asked that a portion of my first week at ICOT be allocated to the completion of that paper. Dr. Furukawa set up a schedule for my first week that included completion of my paper, discussions with various members of the First Laboratory, and a lecture on my representation system DLOG.

2.1 Discussions with First Laboratory researchers

My first meeting was with Dr. Furukawa and Dr. Kunifuji. Dr. Furukawa first provided an overview of the research activities of the first research lab, including the work on inference and problem solving (led by Dr. Kunifuji) and the work on the further development and applicaton of the Guarded Horn Clause language, developed by Kazunori Ueda. Dr. Kunifuji followed with a more detailed presentation of the work on inference and problem solving. I had read two papers on a knowledge assimilation and accomodation system produced by memebers of the first research lab, and I found many similarities in their research and my own, especially my most recent work with Drs. David Poole and Romas Aleliunas, my colleagues at the University of Waterloo. Although I made no firm proposal at this point, collaboration with the inference and problem-solving group seemed most interesting.

A subsequent meeting was arranged between Kazunori Ueda and Akikaza Takeuchi, to discuss the possibility of collaborating on the further development and application of GHC. In discussions with Ueda and Takeuchi, I found that the application of partial evaluation and the development of modules for GHC were among the most active concerns. As I had worked on the development of modules with my graduate students, Mantis Cheng and Chris Baird, and other colleagues at Waterloo, the possibility of cooperating with Ueda and Takeuchi on modules for GHC seemed possible.

At this point, however, I delayed further consideration of a collaborative research topic until after I had the opportunity to lecture about some of my own work. I felt that this would give my ICOT colleagues an opportunity to consider projects that I might be helpful with.

3 Lectures at ICOT

While at ICOT, I gave five lectures:

- DLOG: a logic-based data model for the machine representation of knowledge, Friday, September 6, open lecture, ICOT.
- On overview of research in the Logic Programming and Artificial Intelligence Group at the University of Waterloo, Monday, September 9, First Laboratory, ICOT.
- The D-machine: a simple abstract Prolog machine, Wednesday, September 11, First Laboratory, ICOT.
- Interpreting descriptions in a Prolog-based knowledge representation system, Friday, September 20, ICOT Working Group on Foundations of Artificial Intelligence, ICOT
- What is a knowledge representation system?, invited lecture, Intelligent Knowledge Processing Symposium, Japanese Science Council Lecture Hall, Friday, September 27, Nogizaka.

Material presented at these lectures has been and made available to ICOT members.

3.1 Collaborative research proposal

My New Generation Computing submission was completed on the Tuesday of my second week at ICOT (attached as Appendix A). After completing that paper and two lectures, I began to consider topics for collaborative research. Of particular interest to my colleagues in the First Laboratory was my most recent work with Drs. David Poole and Romas Aleliunas of Waterloo. I had breifly described our work on the theory formation system Theorist, and explained how I believed that it unified much of the research currently in progress in the First Laboratory.

In order to provide some material for concrete discussion, I prepared a proposal of possible collaborative research topics to discuss with Dr. Furukawa. Dr. Furukawa read the proposal, and suggested that I pursue the development of a version of Theorist based on definite clauses. The Theorist program developed with Poole and Aleliunas used the

full clausal form of first order logic as a representation language, and thus required a full clausal theorem-prover as the deductive basis. The initial collaborative research proposal sought to develop a simpler model of Theorist based on definite clauses. Dr. Furukawa pointed out that the classification of definite clauses into facts and integrity constraints provided the necessary concept of potentially inconsistent hypotheses, which is fundamental in Poole's model of theory formation based on deduction. We decided that I would pursue the development of this idea, and consider its application to diagnostic reasoning.

4 Results of collaborative research

During the beginning my third week at ICOT, I summarized the current state of my work on the simplified version of Theorist (henceforth called Theorist-S) and its application in diagnostic reasoning. In particular, I demonstrated a working program that does diagnosis using a definite clause database partitioned into three categories of assertion: facts, integrity constraints, and possible hypotheses. In addition, the structure of the Theorist-S diagnosis system suggested that the form of MYCIN rules was actually an amalgamation of meta and object level information that could be completely separated, and then automatically reassembled using Takeuchi's program for partial evaluation. My brief summary of this work included the abstract and introduction of a proposed paper, by Furukawa and myself. Dr. Furukawa agreed that we should complete the paper (a partial draft and related notes are attached as Appendix B).

In addition to my summary of the Theorist-S program and various comments on its use for diagnosis, I included a list of seven or eight related research topics that arose during the programmng of Theorist-S. I indicated that I would like to discuss these topics with Furukawa and Kunifuji, and perhaps spend the rest of my stay working on one of the proposed related topics. After discussions with Furukawa, Kunifuji, and several other members of the problem solving and inference group, I decided that, in the few days remaining I would devote my time to

- working on the draft of the paper reporting Theorist-S,
- writing this report,
- preparing my lecture for the Intelligent Knowledge Processing Symposium at Nogizaka, and
- working on a generalization of Theorist-S that would consider the use of GHC for the parallel verification of hypotheses in the theory formation computation.

The latter topic seemed to be the most exciting direction to pursue, as it has the potential of unifying the work of Kunifuji's group, the work of the GHC group, and Furukawa's earlier work on Mandala. In particular, there seemed to be an opportunity to exploit parallelism in a semantically useful way, as previously suggested in the original Theorist paper by Poole, Aleliunas and myself.

Most of my final week at ICOT was devoted to considering the possibility of exploiting parallelism in Theorist-S. My initial ideas led me to consider possible process structures for a parallel version of Theorist-S. Further discussion with Dr. Furukawa helped me over my serial "mental block" about GHC. Dr. Furukawa explained how GHC merge networks could be used to synchronize the parallel theory formation processes that I was considering. As time was growing short, I summarized the status of my work by writing an abstract and brief outline for a potential paper on the topic (attached as Appendix C). Further discussions on the parallel Theorist-S continued during the week, with Hirohisa Seki and Akikaza Takeuchi explaining their algorithm for taking two consistent theories and identifying a "crucial literal" whose verification (e.g., by querying the user) would distinguish the two competing theories. Further work was done by Masaru Ohki, who began converting the serial Theorist-S Prolog program to a similar GHC program.

In addition to the work on the parallel version of Theorist, I spend an afternoon with Dr. Kuniaki Mukai discussing the relationship between his CIL programming language and my DLOG system. Both systems make extensive use of complex embedded terms whose manipulation is specified in terms of extended unification. I thoroughly enjoyed this discussion with Dr. Mukai, and learned that we had much in common. In particular, Dr. Mukai verified my intutitions about the relationship between descriptions and the situation semantics of Barwise and Perry, on which Dr. Mukai is an expert. We discussed the general use of embedded terms in logic programming systems and found a great many issues on which our individual research reinforced each others' theories. I promised to send several relevant papers to Dr. Mukai upon my return to Waterloo.

5 Collaboration to be continued

Dr. Furukawa and I decided that we would submit the first paper on Theorist-S to the next International Logic Programming Conference to be held in London next year. I offered to complete the partial draft after returning to Canada, and return it for Dr. Furukawa's comments.

There is substantial work to be done on the parallel GHC version of Theorist. Kazunori Ueda indicated that I could take a copy of GHC with me to Waterloo, in order to continue my attempts at implementing a parellel Theorist program. In addition, I will take a copy of Takeuchi's partial evaluation program. I intend to have some of my students convert Ueda and Takeuchi's programs to run under our WUP version of Prolog, and then continue the work begun at ICOT. I expect that several drafts of a parallel Theorist paper will cross the Pacific before we conclude this initial investigation.

Dr. Mukai and I did not consider any joint research during this visit, but we did promise to keep in close touch about progress in our work, and to consider the possibility of collaboration should I return to ICOT at some future date.

6 The research environment at ICOT

I had some knowledge of the research environment that I would find when I arrived, through discussions with Maarten van Emden who had visited last February. I was, nonetheless, surprised that the open office organization of the laboratories could provide such a congenial atmosphere for productive research. There were numerous informal discussions among researchers in the first laboratory (my limited understanding of Japanese prevented me from participating in most of these). It is important to note that most of these seemed be about research issues—"GHC copy problem" sounds the same in Japanese as it does in English. There were also numerous more organized meetings where some subset of the laboratory would move off to a meeting room, to hold more lengthy discussions. This organization seemed very productive, and I speculate the the well-known Japanese talent for integrating disparate problems and solutions is a result of this kind of atmosphere.

The computing resources of the First Laboratory, and ICOT in general, seemed quite well organized. When I arrived an account on the DEC20 TOPS20 system had already been created, so I could begin my work immediately. Because I came from a VAX/UNIX environment, I initially found TOPS a little clumsy, but I was surprised how little I needed to learn to use the machine effectively.

The ICOT computing facility seems quite well run, but I was surprised how much special knowledge each researcher had to have in order to accomplish the task of preparing a research paper. The current document processing system in use is a TOPS20 version of TEX. I have few kind words for TEX, but most of my criticism is about it's design rather than its effective use at ICOT. The standard TEX output device is a Canon laser printer—the older style, which uses the kerosene-based imaging engine. If you were the first user

of the day, then you had to know how to reset the laser printer, and restart the TOPS20 spooling program. I'm afraid that I have become spoiled by the ease with which I can produce typeset quality printing at Waterloo. Still the system seemed quite well exercised by all the ICOT members, who seem to have no difficulty in performing many aspects of computing operations.

As for the development of logic programs, I found the DEC20 Prolog very stable and productive. However, as the FGCS project has now begun its second phase, it was interesting to find that very few people seem to be using the PSI machines as a personal Prolog development tool. While the research facilities are simliar to the ones that I current enjoy in my own Logic Programming and AI Group laboratory, I believe that some aspects of the FGCS success will depend on delivering the computing power of PSI machines as soon as possible.

7 The FGCS project—prospects for success

There are several ways to judge the success of ICOT and it's ten year FGCS project. The first is from the viewpoint of a non-Japanese observer, who anticipates that success in acheiving the most abstract goals of the FGCS project will not only transform the way in which our civilizations use information, but will have a potentially staggering effect on the computing industry in the rest of the world. This view is the virus that has inflicted the West with "Fifth Generation hysteria," and has precipitated the creation of projects like Britain's Alvey, the EEC's ESPRIT, the United States computer manufacturer's MCC, the Australian Fifth Generation Project, and the Canadian Society for Fifth Generation Research. So far, merely believing that the FGCS project *may be* successful has made great quantities of research money available to Western computer scientists. As I am an academic scientist whose major interests lie in the areas of logic programming and AI, the possibility of a successful FGCS project has contributed to a great improvement in my research environment. Furthermore, since I am an academic, I have no fundamental economic fear related to the success of the FGCS project.

Another view on the potential success of the FGCS project is the scientific view, which requires a careful evaluation of the possibility of achieving success in the stated goal of producing a "knowledge information processing system." There are many fundamental problems to be solved in the last two phases of the FGCS project—the major ones are to develop a simple and efficient way to exploit AND parallelism, and to develop a sound methodology for both that acquisition and delivery of human knowledge. However elusive these goals may seem, it important to acknowledge that much of the problem in achieving these goals lies in the sound integration of the component parts. It is impossible to speculate on the timing of fundamental research breakthroughs, but I feel quite comfortable in suggesting that success in developing a generally useful knowledge information processing system will depend vitally on the integration as well as production of results. In this regard, the Japanese have no equal.

The final way in which I suggest the FGCS project can be viewed is from the viewpoint of the Japanese themselves, especially the Japanese computing research community. From this view, I believe, the project is already a great success. The classical western folk tale of the Japanese lacking in inovation and creative thinking is no where to be found at ICOT. Furthermore, the constant flow of manufacturer's researchers through ICOT has greatly increased the awareness of AI, and is rapidly creating a country that is filled with expertise on everything from parallel architecture to situation semantics. This kind of success is most exciting to me, as I believe that it raises the intellectual level of the general research community which, in turn, fosters better science.

8 Collaboration by visiting researchers

There have been approximately twenty visiting researchers over the first four years of ICOT's existence. I cannot really measure their respective contributions, but I believe that ICOT carefully considers each potential candidate before making an invitation. My only advice on this matter is to encourage the visitor to find out about the research activities at ICOT before hand, and come prepared to explain his own interests as quickly as possible. I think that it is always a mistake to decide a detailed collaborative research plan before you know anything about your collaborators, however ICOT is certainly well-prepared in advance of their visitors, so I would suggest that all potential visitors are likewise prepared.

I might add that I had no problem finding interesting topics to puruse. In fact, quite to the contrary, I found there to be too many interesting things in progress. Here I find that a certain discipline is necessary, especially as the normal duration of collaboration is about three weeks.

9 Conclusions

In case I haven't expressed my positive feelings clearly enough, let me conclude by saying that I have enjoyed this time at ICOT as much as any other research time I have ever had. The place is rich with open-minded clever researchers, who make it a delight to work long hours and pursue that next idea.

Acknowledgements

A major component of successful research is the quality of the human environment. I have had the aid of many kind people, who have done everything from teach me how to use EMACS to point out that I should be careful not to ask for "the edge of a table" when I really want chopsticks. In particular I found Hiroyuki Kusama a splendid host, who seemed to anticipate my every need and desire. Masaru Ohki endured my passion for curry, and contributed his services way beyond what I deserved. Koichi Furukawa was most empathetic about my initial beer allergy, and taught me a lot about doing research besides! I am especially greatful to Ami Senba for teaching me many things about Japan, including all about Japanese baseball.

I am thankful to all the people in the First and Third Laboratories who so graciously endured the strange habits of this foreigner.

Curriculum Vita

Randolph (Randy) George Goebel

August 1985

PERSONAL

Born May 18, 1952, Melville, Saskatchewan, Canada.
Married to Lanie Susann McEachren November 12, 1976.
Two children, Kari Walker born July 7, 1981, Jodi Corine born February 21, 1983.

DEGREES

Bachelor of Science (Computer Science), University of Regina, 1974.

Master of Science (Computing Science), University of Alberta, 1977;
dissertation: *Organizing factual knowledge in a semantic network.*

Doctor of Philosophy (Computer Science), University of British Columbia, 1985;
dissertation: *A logic-based data model for the machine representation of knowledge.*

EMPLOYMENT

Academic

Assistant Professor, Feb85-, Computer Science, University of Waterloo.
Lecturer, Jan82-Jan85, Computer Science, University of Waterloo.
Sessional Lecturer, May79, Computing Science, University of Alberta.
Sessional Lecturer, Jul78-Aug78, Computer Science, University of Regina.
Sessional Lecturer, Jul75-Aug75, Computer Science, University of Regina.

Business

Consultant, Sep83-, Snowbird Computing Consultants Ltd., Waterloo, Ontario.
Consultant, Mar83, IBM Data Systems Division, Poughkeepsie, New York.
Vice-President, Jan76-Aug77, Computek Ltd., Edmonton, Alberta.

INTERESTS

Logic programming, artificial intelligence, knowledge representation systems, logic programming systems, and document preparation systems.

PUBLICATIONS

Refereed

R. Goebel (1985), Interpreting descriptions in a Prolog-based knowledge representation system, *Proceeding of the Ninth International Joint Conference on Artificial Intelligence*, August 16-18, UCLA, Los Angeles, California [to appear].
R.K.S. Lee and R. Goebel (1985), Concurrent Prolog in a multi-process environment, *IEEE 1985*