A report on my visit to ICOT
July 6 - July 24, 1991

Donald Loveland
Duke University

SUMMARY

The main activities of my visit to ICOT were:

- A talk given at LPC'91, on the subject of Disjunctive Logic Programming: Prolog with Case Analysis,
- Discussion with members of ICOT regarding their research,
- A informal discussion of the Model Elimination proof procedure and our recent efforts to extend its power,
- A visit to ASTEM in Kyoto to give a more extensive version of the LPC'91 talk and meet with Dr. Sakama regarding his research in disjunctive logic programming. This was followed by a visit to the Mitsubishi Central Research Laboratory in Osaka also regarding logic programming.

REPORT

I will address each of the topics in the highlights above briefly in this section.

The LPC'91 talk –

Regarding the talk on Disjunctive Logic Programming it seems sensible to include the extended abstract as an appendix to this report rather than reiterate the substance of the talk. I also include further comments on this topic when reviewing the interaction with Dr. Sakama in the last subsection of this Report section. Perhaps an overview remark is reasonable here. I

view the development of the Near-Horn procedure as very successful technically, in that many of the properties that one would wish of a system to extend Horn clause logic programming to the disjunctive domain are realized. (This list of properties was discussed carefully in the talk.) The more difficult problem is whether this whole area is important; does there exist applications for this technology? When I began the investigation I had several small examples of apparent use. As I reflected on applications I discovered many apparent uses of the disjunctive systems could be restated within the Horn clauses domain or even handled directly by use of the alternate answer mechanism, for example. The one area where I see true advantages to this extension I label planning although this label does not always coincide with other people's notion of the word "planning". Planning problems have the characteristic that a single minimal model for the database is inappropriate because several different plausible models exist with later events (i.e., more information) necessary to select the appropriate model. Although we can provide small motivating examples, it remains for an investigator in the planning area to provide a substantial application. When reviewing the talk with me Dr. Hasegawa suggested that there may be an application of this approach to Assumption-based Truth Maintenance systems.

Discussions with ICOT researchers –

Dr. Fujita was kind enough to arrange meetings with various researchers so that I would learn of some of the projects that are being studied in the 5th laboratory. Besides Dr. Fujita himself I met with Drs. Arima, Ohta, Inoue and Kawamura. The work of Dr. Fujita was partially known to me beforehand from meetings in Durham, NC (Duke) and Argonne. The Model Generation Theorem Prover is a very interesting extension of the SATCHMO theorem prover of Manthey and Bry. Technically, it is a hyperresolution system that uses the range restriction condition to assure generation of ground atoms only. (Alternately, the system is a forward chaining system generating ground atoms.) This system is very appropriate for the architecture of the KL1 system on which it is implemented, due to the built-in matching algorithm that suffices when calling goals are ground. Success dealing with a substantial number of problems in terms of performance equal to and usually better than other theorem provers within this domain justifies this approach.

With Dr. Hasagawa I discussed a way to incorporate the idea of relevancy testing into MGTP. The idea is to use MGTP's pure use of forward chaining instead of the backward chaining component used in the paper "Incorporat-

ing Relevancy Testing in SATCHMO" by D.S. Wilson and D.W. Loveland, Duke University Tech. report CS-1989-24, November, 1989. (Dr. Hasagawa has a copy of this paper.) The key point is that the procedure that Debbie Wilson and I found gives a nonobvious search pruning condition that is not a direct extension of Magic sets to nonHorn clause programs. Namely, we show that one can demand that every positive literal of a nonHorn clause be shown relevant before the clause need be called. For MGTP this says that all consequent atoms should be shown relevant before the clause is checked for violation. It is also possible to use only matching for the determination of relevancy although some power is lost by not using unification. However, if matching is much faster, or easier to incorporate into the relevancy procedure for MGTP, then matching may be the correct pattern matching algorithm to use. The key observation here is that relevancy is only a planning device so information can be lost without endangering soundness. The trick is to be sure that information is passed back from goal to called goal but possibly a nonunifying subgoal is labeled relevant that unification would show is not relevant. This means that the pruning is not as strong as it could be with full unification. I think that this "relevancy propagation" could be very useful in some cases for guiding the proof search, and will be very interested in its effect if it is indeed built into MGTP.

Model Elimination presentation –

Dr. Schumann and I gave a joint working group presentation on the Model Elimination(ME) proof procedure and its implementation. My remarks were restricted to the reason for renewed interest in the ME procedure and an outline of the current research in the extension of the procedure. A copy of the report "METEORs: High Performance Theorem Provers using Model Elimination" has been left at ICOT. In summary, the reasons for renewed interest center on the applicability of Prolog implementation technology (the Warren Abstract Machine) because ME is a linear input extension of Prolog to full first-order logic (in essence). The strong implementation techniques apply to sequential and parallel machines. This fact in conjunction with fast workstations and a good implementation by Owen Astrachan that handled well the added burden of the reduction mechanism (for ancestor goal application) allowed two challenge problems of Woody Bledsoe to be solved. No other theorem prover implementing a uniform proof procedure has been able yet to also obtain these proofs. (The STR+VE prover of Hines and Bledsoe do prove these and more but use a special-purpose real number inequality prover among its tools.) Regarding the current research

on ME the focus at Duke is on caching, the use of intermediate results to prevent as much redundant computation as possible. This involves lemmas and "failure lemmas". Lemmas are single-literal clauses (at present) that are logical consequences of the program. Many proofs have the opportunity to use lemmas if they can be effectively restricted in number while retaining the useful lemmas. Also of potential usefulness is the retention of expanded subgoals (which we also call "failure lemmas") which record which goals have been fully expanded and should not be reexpanded in the same environment. The notion of "environment' is complex when the ancestor goals are available for reduction. Appropriate strategies for full use of failure lemmas do not yet exist. A few examples of successful use of lemmas exist but in-depth consideration of caching is just getting underway.

The visit to Kyoto and Osaka –

Both visits focussed on logic programming and included a talk by me on Near-Horn Prolog. However, the two visits were quite dissimilar beyond that point. In Kyoto I was the guest of Dr. Sukama, whose research includes the subarea of disjunctive logic programming. Dr. Sukama's current work pertains to the semantics of Horn clause and disjunctive logic programming. This area has been the focus of substantial research in the last several years because of the need to better understand the meaning of the "closed world assumption" as regards negative queries (goals). Dr. Sakama has proposed a different semantics that can provide for an "inclusive OR" or an "exclusive OR" interpretation within the same semantics. The ambiguity is controlled by the author of the program. The inclusive OR is preferred unless the program author inserts a negative clause denying that interpretation. It is an interesting idea, well conceived and technically well developed. There have been many semantics proposed because none yet seem definitive (and each has its special use depending on the author's intent). I do think that we have hit the point of diminishing returns on the study of semantics for negation, especially since many of the semantics yield difficult (or noncomputable) procedures for determining the negation. Dr. Sakama also feels that the research in the semantics area is concluding. He has a good understanding of the general picture and I do not expect him to continue in this area when the interesting questions are closed.

The audience at ASTEM consisted of four people, two of whom could follow me very well, so I spoke quickly and covered the near-Horn procedure, the fixpoint semantics associated with the procedure and the consequence of adding closed-world-assumption-style negation to the procedure. At Mit-

subishi I spoke to logic programmers not familiar with the idea of disjunctive logic programming so I presented only the procedure. The group was reasonably large (about ten people) but interaction was very good so I was able to present the procedural aspect in depth and feel that many understood both the procedure and the reason for the choice of that design.

Two researchers at Mitsubishi presented their current work and I found the work quite interesting and current. Dr. Seki has developed and implemented an extended form of magic sets for deductive databases. Dr. Takahashi has developed a language for AND-OR parallelism that is implemented in KL1.

I enjoyed both visits and particularly found the interaction with Dr. Sakama worthwhile. I thank ICOT for the opportunity to make these visits.

ACKNOWLEDGMENTS

I wish to extend many thanks to the members of ICOT, and others, that have made this visit very worthwhile for me. I thank the program committee of LPC'91 and especially Dr. Amamiya and Dr. Aiba for extending the invitation to speak at LPC'91. For the invitation to attend ICOT after the LPC'91 I thank Dr. Fuchi. Dr. Furukawa and Dr. Hasegawa were very kind to make my stay in their laboratory as pleasant and educational as it was. Special thanks go to Drs. Fujita and Iwata who spent much time seeing that my stay at ICOT went smoothly. The visit was also made much more pleasurable by the invitations to visit the homes of the Fuchis and Furukawas and to enjoy very fine Japanese meals and wonderful hospitality. There are many people who have contributed in many ways and I wish to thank them all collectively. It has been of great interest to me to see the physical environment, compact by U.S. standards, which has been so productive of ideas and influence in the fields of logic programming and parallel computation, and to see the rapid increase in sophistication and effectiveness that the research in automated theorem proving and program synthesis application is achieving.

CURRICULUM VITAE

Donald Loveland

May 1991

PERSONAL DATA

Born: December 26, 1934, Rochester, New York

Married, two children.

EDUCATION

A.B., Oberlin College, 1956

S.M., Massachusetts Institute of Technology, 1958

Ph.D., New York University, 1964.

EXPERIENCE

Mathematician Programmer, IBM, Yorktown Heights (1958-1959);

Research Assistant, New York University (1960-1961);

NSF Fellow, New York University (1961-1963);

Instructor, New York University (1963-1964);

Assistant Professor, New York University (1964-1967);

Assistant Professor of Mathematics and Computer Science, Carnegie-Mellon University (1967-1969);

Associate Professor (1969-1973);

Professor and Chairman, Computer Science Department, Duke University (1973-1978).

IBM Distinguished Faculty Visitor
IBM T.J. Watson Research Center, 9 mos. (1979-1980)

Professor, Computer Science Department, Duke University (1978-)