AD-A162 762 INTELLIGENT AGENTS PROJECT(U) STANFORD UNIV CA DEPT OF COMPUTER SCIENCE 03 DEC 85 N00014-81-K-0303 1/1

UNCLASSIFIED F/G 9/2 NL

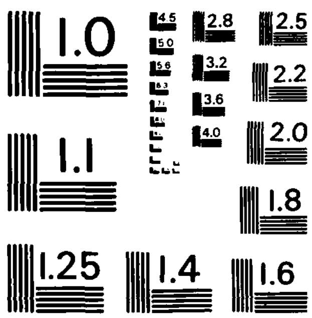

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

Stanford University
Knowledge Systems Laboratory
Department of Computer Science
Pala Alto, CA 94304

N00014-81-K-0303

12

DTIC ELECTE DEC 30 1985 S A

Final Report
Intelligent Agents Project

December 3, 1985

1 Introduction

As suggested in the reporting guidelines for ONR information systems program contractors, this report summarizes the research progress made at Stanford University under ONR Order No. NR 049-489/8-6-82, Contract No. N00014-81-K-0303. The overall purpose of the contract was to support basic and applied research in the study of interacting "intelligent agents" (IA's), each of which was capable of acting autonomously in a precisely specified domain. The principal application domain was that of an intelligent interface to a computer operating system.

Work in this area split naturally into two broad subareas. The first involved construction of an individual intelligent agent, while the second required investigation of the issues involved in the interaction between a variety of such agents. Keywords: Information Systems program contractors

The construction of a single autonomous agent in a complex domain itself involved a variety of problems. Firstly, planning research needed to be done to investigate the problems that will be encountered by a planner that needs both to observe its environment before making plans and to execute these conditionally, allowing for the possibility of failure along the way. Intelligent interaction with computer operating systems, including access to remote devices, requires that both of these problems be addressed.

Given these constraints, the domain is sufficiently complex that "blind" planning and inference are unlikely to manage the combinatorial difficulties of the problems encountered. Research was therefore required on the control of such inference, as the IA needs to be able to introspect and control its own activities.

Finally, an individual IA will need to present its results to the user. This led us to investigate the questions involved in intelligent presentation, as well as those in intelligent planning and controlled inference.

Interaction between IA's is a related but separate issue. Research in this area has had to address questions involving general problems of interaction between agents whose goals or world models may conflict; these questions had gone nearly untreated in the AI literature until the IA project addressed them.

2 Sensory and Conditional Planning

The Intelligent Agents domain is one in which many planning problems do not have guaranteed linear solutions. Instead, it is necessary to take steps to gather information about the state of the world and to then act upon the information gathered. This gives rise to plans which, instead of being a simple sequence of actions, contain conditional IF-THEN-ELSE constructs. In addition, such plans may also contain "sensory" actions, the purpose of which are to gather information rather than to change the state of the world. An approach to planning has been developed which allows the construction of conditional-sensory plans in a natural manner, including cases where the sensory actions themselves have prerequisites which must be achieved [3]. The formalism developed for planning has proven useful in other design problems [7,8,9], and has also given rise to research on efficiency gains made by exploiting existing design constraints [4,5,6]. Finally, since most planning and error recovery in the IA domain is fairly standardized, an "expert system" style planner was developed with extensive abilities to transfer files between numerous sites over three networks and to recover from a large number of commonly occurring errors [1,2].

3 Control of Inference

As remarked in the introduction, the need to use more than "canned plans" in the IA domain led to a clear need for better and more general strategies for control of inference than had previously been available. The IA domain poses many problems in which the are large conjunctive queries to solve. If improperly ordered, these problems are often intractable, but use of knowledge about the number of answers for the individual conjuncts can lead to dramatic decreases in the number of possible answers which must be examined. A number of results were obtained on this problem [26,29]. Another area of control research investigated the problem of deciding when a search could be stopped because all possible answers to the original query had already been found [26,27,28].

This work dovetails neatly with the more general work on control of inference which is also in progress at Stanford. An area of this more general control work which was investigated under the IA project involved studying the tradeoff between devoting resources in an attempt to find the best course of action in solving a problem and devoting those same resources to a direct attempt at a solution. A variety of situations were studied and a number of "break-even" points were discovered [19].

4 Intelligent Presentation

An IA's activity will frequently involve the presentation to a human user of informaion about the state of the IA's world. If, for example, a plan has failed unexpectedly, it is far better that the user be presented with a sketch of what has happened and where the world stands now rather than that the machine simply aborts the plan in some unknown state. Not surprisingly, however, the IA will have far more information about the state of the world than should be presented to the user, and the IA will therefore need to present only a suitable subset of the information available to it.

Flexibility in the nature of the information being presented requires flexibility in the method of presentation as well, since different sorts of facts should generally be presented using different representation methods. An investigation of these issues led us to explore a variety of criteria for automatic generation and evaluation of methods for machine presentation of information [12,13,14,15].

5 Interacting Agents

The IA includes many "canonical agents" working together. However, it is not always possible or even desirable for the individual agents to blindly obey one another; to lay the foundation for an investigation of the details of multi-agent interaction, we began by defining notions of rationality for single agents. The first case investigated dealt with agents in a master/slave relation, so that each agent was able to fully specify the actions of other agents in order to carry out a task. A solution to the problem of ordering activities among agents was discovered [16].

The next problem considered dealt with the discovery of communication strategies between agents with identical goals but with incomplete or conflicting world models [17]. Following this, we were finally able to turn to the case of fully autonomous agents that would not necessarily cooperate but might instead pursue courses of action intended to achieve potentially conflicting goals. A hierarchy of rationality assumptions was developed, and we investigated the consequences of each agent's assuming that the behaviors of the others could be described by one or more of the definitions in the hierarchy. The formalism developed allowed for the modelling of restrictions on communication and the exchange of binding promises among agents. The work is described in [10,11,18,20]

Another area of research was on the tradeoff between communication costs and parallelism. It was discovered that the communication costs incurred in a joint effort may outweigh the advantages of parallel processing in some situations, and an efficient communication protocol called ESP was developed to address these problems. We also investigated a "Variable Supply Model," which covers a large spectrum of strategies in the tradeoff between parallelism and communication. A case study involving run-time allocation of deductions to multiple agents is presented in [24] and shows various empirical breakeven points in the trade between parallelism and reduced communication costs. Another case study is presented in [22] and investigates the static allocation of deductions to a large number of agents. This study attempts to make compile-time estimations of run-time communication costs between the agents and, based upon these estimations, to distribute the database among the various agents such that agents with a great need to communicate will be near to one another. In addition, the system can fold parts of the database together onto the same agent in an attempt to reduce communication costs at the expense of reduced parallelism. Relevant references are [21,22,24,25,23].

References

[1] Jonathan Cronin. The intelligent agent: an expert system. June 1984. Master's Practicum.

[2] Michael Dolbec. A planner for the intelligent agent. June 1982. Master's Practicum.

[3] J. J. Finger. *Sensory Planning.* Technical Report HPP-82-12, Stanford University, April 1982.

[4] J. J. Finger and Michael R. Genesereth. *RESIDUE - A Deductive Approach to Design.* Technical Report HPP-83-46 working paper, Stanford University, December 1983.

[5] J. J. Finger and Michael R. Genesereth. *RESIDUE - A Deductive Approach to Design Synthesis.* Technical Report HPP-85-1, Stanford University, January 1985.

[6] Joseph Jeffrey Finger. *Exploiting Design Constraints.* PhD thesis, Stanford University, 1986. (in preparation).

[7] ed. Genesereth, Michael R. *The MRS Casebook.* Technical Report HPP-83-26, Stanford University, May 1983.

[8] Michael Genesereth. *Partial Programs.* Heuristic Programming Project Memo HPP-84-2, Stanford University, November 1984.

[9] Michael R. Genesereth. The use of design descriptions in automated diagnosis. *Artificial Intelligence*, 24:411–436, 1984.

[10] Michael R. Genesereth, Matthew L. Ginsberg, and Jeffrey S. Rosenschein. *Cooperation without Communication.* HPP Report 84-36, Heuristic Programming Project, Computer Science Department, Stanford University, September 1984.

[11] Michael R. Genesereth, Matthew L. Ginsberg, and Jeffrey S. Rosenschein. *Solving the Prisoner's Dilemma.* HPP Report 84-41, Heuristic Programming Project, Computer Science Department, Stanford University, November 1984.

[12] Jock Mackinlay. *Intelligent Presentation: The Generation Problem for User Interfaces.* Heuristic Programming Project Memo HPP-83-34, Stanford University, March 1983.

[13] Jock Mackinlay and Michael R. Genesereth. Expressiveness of languages. In *Proceedings of the National Conference on Artificial Intelligence*, Austin, Texas, August 1984.

[14] Jock D. Mackinlay. *Automatic Design of Graphical Presentations.* PhD thesis, Stanford University, 1986. (in preparation).

[15] Jock D. Mackinlay and Michael R. Genesereth. Expressiveness and language choice. *Data & Knowledge Engineering*, 1:17–29, 1985.

[16] Jeffrey S. Rosenschein. Synchronization of multi-agent plans. In *Proceedings of The National Conference on Artificial Intelligence*, pages 115–119, The American Association for Artificial Intelligence, Pittsburgh, Pennsylvania, August 1982.

[17] Jeffrey S. Rosenschein and Michael R. Genesereth. *Communication and Cooperation.* HPP Report 84-5, Heuristic Programming Project, Computer Science Department, Stanford University, March 1984.

[18] Jeffrey S. Rosenschein and Michael R. Genesereth. Deals among rational agents. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 91–99, The International Joint Conferences on Artificial Intelligence, Los Angeles, California, August 1985. Also published as HPP Report 84-44, Heuristic Programming Project, Computer Science Department, Stanford University, December 1984.

[19] Jeffrey S. Rosenschein and Vineet Singh. *The Utility of Meta-level Effort.* Report No. HPP-83-20, Heuristic Programming Project, Computer Science Department, Stanford University, March 1983.

[20] Jeffrey Solomon Rosenschein. *Rational Interaction: Cooperation Among Intelligent Agents.* PhD thesis, Stanford University, October 1985.

[21] Vineet Singh. *Distributing Deductions to Multiple Processors.* PhD thesis, Stanford University, June 1986.

[22] Vineet Singh and Michael R. Genesereth. *PM: A Parallel Execution Model for Backward-Chaining Deductions.* KSL Report KSL-85-18, Stanford University, May 1985. Submitted to Future Computing Systems.

[23] Vineet Singh and Michael R. Genesereth. *A Variable Supply Model for Distributing Deductions.* Technical Report HPP-84-14, Heuristic Programming Project, Computer Science Department, Stanford University, May 1984.

[24] Vineet Singh and Michael R. Genesereth. A variable supply model for distributing deductions. In *Proceedings of IJCAI-85*, Morgan Kaufmann Publishers Inc., August 1985. Submitted to Journal of Parallel and Distributed Computing.

[25] Vineet Singh and Michael R. Genesereth. A variable supply model for distributing deductions. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 39–45, The International Joint Conferences on Artificial Intelligence, Los Angeles, California, August 1985.

[26] D. E. Smith. *Controlling Inference.* PhD thesis, Stanford University, August 1985.

[27] D. E. Smith. Finding all of the solutions to a problem. In *Proceedings of the National Conference on Artificial Intelligence*, American Association for Artificial Intelligence, August 1983.

[28] D. E. Smith and M. R. Genesereth. *Controlling Recursive Inference.* Heuristic Programming Project Memo HPP-84-6, Stanford University, 1984.

[29] D. E. Smith and M. R. Genesereth. Ordering conjunctive queries. *Artificial Intelligence*, 25, 1985.

END

FILMED

2-86

DTIC