Thesis Proposal: Analysis of Single Function Agents and their Negotiation Behaviour in Expert Systems

Bertram V. Dunskus
Computer Science Department
Worcester Polytechnic Institute

Advisor: Prof. David C. Brown

February 23, 1994

Abstract

Design systems can be built from multiple expert systems called agents. A Single Function Agent is an agent, with only one function, one point of view, and one target object to act on. This thesis will investigate the possibilities and implications of this concept, analyzing the language, negotiation and strategies available to the agents. We hypothesize that negotiation between the agents becomes simpler than it is with agents of larger size. Negotiation knowledge might even be stored in a pattern oriented representation, speeding up the problem-solving process and simplifying the design.

1 Introduction

As maintenance, code reuse and flexibility become increasingly important, current developements in expert systems show a trend towards using multiple agents working on multiple tasks instead of monolithic inference engines. Multi-Agent systems promise to offer reliability even with inconsistent or uncertain data and knowledge. In addition to that, maintenance becomes easier, as the information pertinent to a specific type of expertise is stored in a distinct agent. Also, it is possible to build a library of agents with different domain knowledge and different tasks, and to combine them to form expert systems for different problems.

Conflicts will naturally arise in this kind of situation, where multiple agents with different knowledge have opposing views about the same object. The resolution process should be kept general so that it is not too different for each pair of conflicting agents. Negotiation can be used to resolve such conflicts in the problem solving process.

2 Background

2.1 Multi Agent Expert Systems

Several systems have used a multi agent design. They differ in the size of the agents used, the availability of a central control agent, the integration of external software modules and the amount of collaboration and negotiation involved in the problem solving process.

Some of the research has concentrated on systems of independent agents which are built to have no negotiation [Victor et al. 93], e.g. only one agent has the right to decide on a certain design value. Most of the research concerns systems made from fairly complex agents, with multiple things to decide, having large knowledge bases and complicated interaction structures [Klein 91] [Sycara 90].

Another issue in the design of multi agent systems is the problem of communication, i.e. how information is transferred between the agents. Some of the possible approaches are blackboards, messages and function calling, while the content is often formed using speech acts [Bussmann & Mueller 93] [Taleb-Bendiab & Oh 93].

An approach to standardize the content of the messages, by building common ontologies, as well as to supply a set of communication channels and message wrappers, has been a continuous effort by researchers at Stanford [Kuokka et al. 93]. They have developed a standard for logical expressions [Genesereth et al. 92] as well as a common language for agent communication [Finin et al. 93]. Both of these have been succesfully applied to several projects.

2.2 Negotiation

Sycara states that the requirements for models of negotiation are as follows [Sycara 90]: iterative rather than one-shot; reception of and evaluation of feedback about propositions; proposal of suitable modifications; prediction of possible improvements through proposals; justification and arguments. Her model is built around agents that use both case-based reasoning and utility-based negotiation. The agents store their goals and their expertise in qualitative dependency graphs. When a conflict arises, they try to find a solution that is acceptable to both agents, but still improves their own evaluation of the design. The communication vocabulary is implicitely defined by the overlapping expertise between the two agents.

Sycara's model has the disadvantage that the kind of knowledge the agents have to store is very complex. The conflict resolution needs a lot of qualitative reasoning. In addition to that, one of her basic assumptions is that the agents share common goals, which doesn't have to be the case. Also, the language in which the agents communicate is limited in its expressiveness.

Klein concentrates on the necessary steps to find the best way to solve a conflict [Klein 91]. He uses a hierarchical tree which stores conflict situations at different levels of abstraction. The agents search this tree, until a matching description is found, then they look in a similar tree for solutions. In a sense, the system solves conflicts without negotiation, because all the agents have the same conflict resolution (CR) knowledge and they automatically agree on a CR plan, if one is found.

2.3 Single Function Agents

At WPI three systems using Single Function Agents (SiFAs) have been developed, so far:

- I3D, a system for integrated part design and manufacturing plan production for Powder Processing Applications, used a set of cooperative agents. They acted as a team of experts, assisting the human designer as s/he designed simple powder ceramic components. During both conceptual and detailed design, the agents offered cost estimation, material selection, process simulation and inspection planning. The 16 agents were sequenced, with possible conflicts anticipated and removed at development time.

- I3D+ addressed some of the limitations on conflicts and sequencing in the I3D system. It introduced an agenda-based scheduling system that kept track of agents competing for execution. It also allowed multiple agents to have conflicts about values of the design object. They would negotiate with one another, following a different scheme depending on which of seven local/global goal situations the agents were in.

- SNEAKERS was a system intended for training users in the field of Concurrent Engineering. The user interacts with a set of agents that have different points of view (design, manufacturing, assembly, cost, marketing etc.) and different functions (advising, criticizing, suggesting, analyzing and evaluating). As the user makes decisions to build a model tower, the agents give their respective comments, evaluations and recommendations about the user's design choices.

What these systems have in common is a structure based on a set of small agents, each with an individual, well defined *role* and a *target* object to act on. In I3D+ this concept was extended through the notion of a *point-of-view*, thus defining agents from three sides, e.g. an agent might be a critic (role) of material (target) from the point of view of cost.

3 Proposed Approach

3.1 Goals

The thesis will explore the Single Function Agent paradigm in a more thorough manner than has been done in previous research. After defining a domain independent set of agents we will investigate negotiation, analyzing which pairs/groups of agents would have reason to communicate, and what the information passed between them should be. Then we will investigate the patterns of messages that will appear and the kinds of knowledge that are needed for the negotiation. Knowledge representation will be an issue here. If possible, a catalog of conflicts for Single Function agents, with attached methods of resolution will be developed. Finally, we will look at how much a negotiation history would be of help to the agents.

In order to apply the gained insights to a real life domain, we will build and experiment with a prototype system to design the same object as the I3D system. This will allow a comparison, and it will make the acquisition of specific design knowledge unnecessary, as it is already available.

3.2 Phases

The research will be carried out in phases. As the outcome of later research will give new insights to topics investigated in the beginning, an iterative element will probably appear in the work.

1. Design of a standardized agent structure:
   The first work will define a structure on which all the SiFAs will be built. This will include a first cut at the needed communication primitives and the necessary data storage.

2. Implementation of a communication language:
   In order for the agents to exchange knowledge and data, a standard language, common to all agents will be needed.

3. Implementation of the I3D+ conflicts:
   The conflicts investigated in the I3D+ work will be the first to be implemented. They will form the foundation for further exploration.

4. Analysis and optimization of the negotiation behaviour:
   This will include: gathering knowledge of the best forms for storage of negotiation knowledge, finding the optimal sequences of knowledge exchange, and analyzing what pairs of agents make the most use of negotiation.

5. Summary of the observed performance and patterns:
   A complete description of conflict types and their solution processes for SiFAs, as well as their effect on system performance will be compiled. This will help in deciding the usefulness of SiFAs in future systems and it will allow constructive advice to be given when building one.

3.3 Evaluation

Because the thesis strives to extend and solidify previous ideas about SiFAs, it will be fairly difficult to evaluate the work as success or failure. However, the following criteria can help when evaluating the outcome:

- Comparison of the design objects:
  How good is the quality of the object designed by the negotiating system, as compared to non-negotiating systems?

- Expert Opinion:
  Does the agent negotiation make sense to a human, i.e., is the process understandible at run-time, would a human agree on the outcome, and could one easily read the knowledge representation?

- Performance:
  What is the overhead incurred through negotiation? How does the amount of conflict affect the run-time?

- Comparison to other negotiating systems:
  What is the performance of a system based on SiFAs, compared to a system with a larger agent size?

If these questions can be answered by the thesis, it should be considered as having achieved its goals.

4 Deliverables

- Implementation of a Communication Language:
  A language for interagent communication for SiFAs, probably based on the KIF and KQML standards [Genesereth et al. 92] [Finin et al. 93].

- Hierarchy of Conflict Situations:
  A compilation of all conflict situations open to negotiation, in hierarchical fashion, similar to Klein's work [Klein 91], but specialized for SiFAs.

- Suggestions for future developement and implementation of SiFAs:
  Supportive information for applications and research.

5 Resources

The following combination of hardware and software will be used:

- Hardware: SUN SPARC IPC or DEC Alpha 3000 in WPI's AI lab.

- Expert System Shell: CLIPS C-Language Integrated Production System, with Dyna-CLIPS extension, if available on time for CLIPS 6.0.

- Programming Language: GNU C or GNU C++

References

[Brown 93] D. C. Brown. *Exploring the Role of Single Function Agents in Negotiation.* AAAI Workshop on Cooperative Design, 1993

[Bussmann & Mueller 93] S. Bussmann & J. Mueller. *Bargaining Agents.* Submitted to IJCAI, 1993

[Douglas et al. 93] R. E. Douglas, Jr., D. C. Brown & D. C. Zenger. *A Concurrent Engineering Demonstatration System for use in Training Engineers and Managers.* Applications of AI in Engineering VIII. Vol 1: Design, Methods and Techniques. Proc. AIENG '93, Toulouse, France. (Eds.) G. Rzevski, J. Pastor & R. A. Adey, Computational Mechanics Publications & Elsevier Applied Science, pp. 505-520.

[Genesereth et al. 92] M. R. Genesereth & R. E. Fikes. *Knowledge Interchange Format Version 3.0 Reference Manual* Computer Science Dept., Stanford University, CA 94305, 1992

[Klein 91] Mark Klein. *Supporting Conflict Resolution in Cooperative Design Systems.* IEEE Transactions on Systems, Man, and Cybernetics, Vol. 21, No. 6, November/December 1991

[Finin et al. 93] T. Finin, J. Weber, G. Wiederhold, M. Genesereth, R. Fritzon, D. McKay, J. McGuire, R. Pelavin, S. Shapiro & C. Beck. *DRAFT Specification of the KQML Agent-Communication Language* The DARPA Knowledge Sharing Initiative External Interfaces Working Group, 1993

[Kuokka et al. 93] D. R. Kuokka, J. G. McGuire, R. N. Pelavin, J. C. Weber, J. M. Tenenbaum, T. Gruber & G. Olsen. *SHADE: technology for knowledge-based collaborative engineering*, AAAI Workshop on Collaborative Design, 1993

[Sycara 90] Katia P. Sycara. *Cooperative Negotiation in Concurrent Engineering Design.* Cooperative Engineering Design, Springer Verlag Publications, 1990

[Taleb-Bendiab & Oh 93] A. Taleb-Bendiab & Vincent Oh. *Speech-Act based Communication Protocol to support Multi-Agent Cooperative Design Systems.* Proceedings of the 1993 AI in Engineering Conference, p. 107

[Victor et al. 93] S. K. Victor & D. C. Brown et al. *Using Multiple Expert Systems with Distinct Roles in a Concurrent Engineering System for Powder Ceramic Components.* The Center for Intelligent Processing of Materials, Worcester Polytechnic Institute, 1993