USING SINGLE FUNCTION AGENTS TO INVESTIGATE NEGOTIATION

David C. Brown, Bertram Dunskus & Dan Grecu

Computer Science Dept., Worcester Polytechnic Institute, Worcester, MA 01609

Phone: (508) 831-5618
Email: dcb@cs.wpi.edu

Keywords: *Design, Single Function Agents, Negotiation, Learning*

1. Previous Work

The initial motivation for our research was the desire to build working Concurrent Engineering (CE) systems. Within the last two years we have built two systems with many CE characteristics. As it is difficult to produce a precise description of CE on which all researchers would agree, we will refrain from claiming that they *are* CE systems. However, most would agree that the essential quality is that at design time it is important to consider the impact of the design on aspects from later in the life-cycle, such as packaging, or recycling. The systems built support these considerations.

One of these systems, the I3D system [Bausch et al 1992] [Victor et al 1993], used multiple, cooperative, intelligent agents to serve on a design team with a human designer. Together they designed simple powder ceramic components, using a conceptual design phase and a detailed design phase. The agents were able to respond and assist in both of the design phases, offering cost estimation, material selection, process simulation and inspection planning, for example. There were 16 intelligent agents in I3D.

The second system, called SNEAKERS [Douglas & Brown 1993], was intended to be used during training courses to help educate users about the importance and power of CE. It is a single-user system, with the other members of the CE team being simulated by expert systems (agents). The user designs a tower on the screen, according to some requirements, by selecting from, and specializing, parts offered by the system. As in I3D, the agents offer a variety of comments which the user can act on or ignore.

These systems have several things in common. Those relevant to this workshop are: they used intelligent agents, with each agent having a single function; the agents are only allowed to activate in a fixed sequence.

However, in these systems all possibility of conflict, and hence negotiation, was compiled out by careful ordering of the agents and by careful grouping of the decisions being made. In a new version of I3D (I3D+), to address grouping, we have separated out the decisions into different agents, and, to address ordering, have no predefined order of execution.

2. Single Function Agents

In I3D+ and SNEAKERS we experimented with single function agents. Agents were only able to perform one of a few simple roles [Brown 1992]. The roles used included providing Advice, Analysis, Criticism, Estimation, Evaluation, Planning, Selection and Suggestion.

These roles can be contrasted with the target about which they can make comments. For I3D+ these include Material, Process, Manufacturing, Inspection, Cost, Reliability and Durability. The roles and targets define a 2D matrix of possible agents, not all of which are sensible.

In I3D+ we refined this further by considering another axis, point-of-view, thus classifying agents using a 3D matrix (i.e., role, target, point-of-view). For example, one agent might produce criticism (role) of material (the target) from the point-of-view of cost. Note that, as before, not every combination is sensible. However, this system allows us to quite precisely define every agent in the system. It will also force us to be very explicit about what each agent can do, and what it knows, while allowing us freedom to implement the agents in any way appropriate, provided they play their appointed role. We refer to these limited agents as *Single Function Agents*, or SiFAs.

We consider the functionality to be more general than target and point-of-view. The function defines the agent, and the two other aspects are instantiating it at the domain level. As we experiment with negotiating SiFAs, it is possible that we will discover that new functions will be required, or that some of our current list of functions should be split or combined. We would like to discover the most appropriate level of granularity for design systems involving negotiation.

We expect to learn a lot about the roles of different types of knowledge in different negotiation schemes. We have already added Praisers to the list of types of roles. Where grouping of decisions allowed us to avoid negotiation in I3D and SNEAKERS, this approach of strict separation of role will force us to face negotiation directly and in detail.

For example, it is already clear that Critics and Estimators interact in interesting ways with the agents doing design (Advisors), allowing several ways to negotiate between them. For example, a critic can criticize an estimate: what do they negotiate about? Perhaps the method of estimation? Or the data used? Can Critics conflict? What if Praisers and Critics disagree?

This research will investigate the possibilities and implications of the SiFA concept, analyzing the communication language, negotiation and strategies available to the agents. We hypothesize that negotiation between the agents becomes simpler with SiFAs than it is with agents of larger size. This is because SiFAs heavily constrain the possible types and topics of negotiation. Negotiation knowledge might even be stored in a pattern oriented representation, speeding up the problem-solving process and simplifying the design. The resulting "dissection" of negotiation should more explicitly reveal the trade-offs being made during the design.

We also expect to discover generic patterns of use of functions in the negotiation process. For example, it may be that certain types of conflicts, and therefore certain types of negotiation, are always associated with the same type of cluster of SiFAs. This sort of observation about the nature of negotiation between agents would only be possible with functions at this level of granularity.

We expect to be able to show by using templates, and an object oriented implementation using inheritance, that SiFAs of similar type share characteristics. This should lead to ease of development of SiFA systems, some reuse of SiFAs, and the possibility of using small Knowledge Acquirers tuned to each type of SiFA, and perhaps to types of conflicts and negotiation knowledge.

3. Current Research Directions

Current research is proceeding in two directions. The first is to study SiFA systems to see what kind of communication and negotiation arises naturally, and build implementations to explore this. The second direction, longer term, is to study how SiFA agents might learn by exchanging goals and knowledge.

3.1 A Study of Negotiation in SiFAs

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.

Because the research 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 understandable 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? Where is that overhead with respect to types of SiFAs? 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?

The expected outcomes from this short-term work are:

Implementation of a Communication Language: A language for interagent communication for SiFAs, probably based on the KIF and KQML standards [Genesereth & Fikes 1992] [Finin et al 1993].

Hierarchy of Conflict Situations: A compilation of all conflict situations open to negotiation, in hierarchical fashion, similar to Klein's work [Klein 1991], but specialized for SiFAs.

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

3.2 Learning with SiFAs

Single function agents are supposed to allow the definition, implementation, testing and evaluation of elementary patterns of conflict, communication and negotiation. In order to generate such patterns agents are classified according to their functionality, domain and point of view. Assuming that the possible patterns of conflict are established, one first step further would be to explore what the agents can share along each of these dimensions. This is important in order to determine what an agent can understand about another agent. Functionality sharing can mean understanding of design or critiquing strategies, a common domain would possibly allow for exchange of domain rationale, while common points of view are based on similar goals.

This research will attempt to lay the groundwork for a framework for learning during, and due to, SiFA negotiation. We are considering the following possibilities:

- Knowledge about another agent allows one agent to build a model of the other agent's specific intentions and beliefs. In the case of SiFAs, such models can be easily kept within reduced dimensions, assuming that the classification of the other agent is known.

• Understanding the other agent's domain/goal/functionality can help the agent define expectations about the other agent.

• This understanding helps anticipate possible conflicts between agents. Reasoning to this purpose can result in the compilation of negotiation knowledge.

• Negotiating with different types of agents could generate learning about how to adequately select/construct negotiation strategies. Success of negotiation is highly dependent on finding a way to negotiate which is suitable for both parties involved.

• Negotiating is a rich source of learning how to relate intentions of another agent to his functions, goals (and, if known, beliefs).

• The SiFA model is supposed to offer a clear overview of the sources of conflict and of the negotiation. Extended explanation facilities about the conflict and the negotiation process can be implemented. Agents should also be sensitive to learning generic patterns of conflict generating sources.

The necessary exchange of knowledge between the agents to achieve these goals does not conflict with their restricted functionality. The negotiation-related communication is assumed to take place at a meta-knowledge level and does not contradict with their domain functionality.

Further potential advantages of SiFA architectures would be the development of "pure" utility functions. Such functions would reflect well-defined positions in the agent space and, together with a combination method, would offer a more accurate determination of the negotiation result.

Single-function agents research is eventually supposed to provide an insight into efficient function, domain or view combinations in the attempt to define complex agents. Superimposing single-function agents (possibly through coalition formation) might be a relevant method to produce such agents. We hypothesize that patterns of negotiation will help identify which agents could be put together to form complex agents. It may also be possible to identify types of complex agents dependent on which SiFAs they contain.

4. Conclusion

Despite some very interesting work on computational models of negotiation -- such as that by Sycara [1990], Lander & Lesser [1992], Klein [1991], Werkman [1992], Kannapan & Marshek [1992], and others -- we feel that there is still much to be done. Despite having contacted many people directly and via email, one surprise is that there appear to be very few implementations, few computational models, little evaluation of the systems, and no comparative evaluation.

We will test our hypotheses about the expected advantages of SiFAs by building design systems in several areas, and will try to compare them with non-negotiating versions of the systems.

We consider that a methodological disadvantage of complex agents is that they can span several domains/functions/points of view. The possibility of classifying agents according to these criteria is highly improbable and therefore learning to associate intentions and beliefs with functions and goals may be less straightforward than in the SiFA case.

In the negotiation process an agent has to:

a) acquire knowledge describing and defining the other agent (part of which is implicit in the single function case assuming the type of the other agent is known);

b) learn how to react to possible patterns of intentions/beliefs/functions/goals and form expectations;

c) abstract general negotiation answers and compile negotiation strategies.

We feel that it will be profitable to proceed with knowledge acquired during negotiation with single function agents and then to compose elementary negotiation strategies for complex agents, rather than by starting the learning process directly with complex agents.

5. References

J.J.Bausch, D.C.Zenger, D.C.Brown, R.Ludwig & R.D.Sisson (1992) Integrated Design and Manufacturing Strategies for Powder Processing Applications. Proceedings, University Programs in Computer Aided Engineering Design and Manufacturing (UPCAEDM'92), Tennessee Technological University.

D.C.Brown (1992) Design. Encyclopedia of Artificial Intelligence, 2nd edn., S.C.Shapiro (Ed.), J.Wiley.

R.E.Douglas & D.C.Brown (1993) A Concurrent Engineering Demonstration System for use in Training Engineers and Managers. Proc. AI in Engineering Conference.

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

M.R.Genesereth & R.E.Fikes (1992) Knowledge Interchange Format Version 3.0 Reference Manual, Computer Science Dept., Stanford University, CA 94305.

S.M.Kannapan & K.M.Marshek (July 1992) A Schema for Negotiation between Intelligent Design Agents in Concurrent Engineering. In: D.C.Brown, M.Waldron & H.Yoshikawa (Eds.), Intelligent Computer Aided Design, Elsevier Science Publishers (North-Holland).

M.Klein (1991) Supporting conflict resolution in cooperative design systems. IEEE Trans. Systems, Man, and Cybernetics, November/December, Vol. 21, No. 6, pp. 1379-1390.

S.E.Lander & V.R.Lesser (1992) Customizing Distributed Search among Agents with Heterogeneous Knowledge. Proc. 5th Int. Symp. on AI Applications in Manuf. & Robotics, Cancun, Mexico.

K.P.Sycara (1990) Cooperative Negotiation in Concurrent Design. In: Cooperative Engineering Design. Springer Verlag.

S.K.Victor, D.C.Brown, J.J.Bausch, D.C.Zenger, R.Ludwig & R.D.Sisson (1993) Using Multiple Expert Systems with Distinct Roles in a Concurrent Engineering System for Powder Ceramic Components. Proc. AI in Engineering Conference.

K.J.Werkman (1992) Multiple Agent Cooperative Design Evaluation using Negotiation. In: J.S.Gero (Ed.), Artificial Intelligence in Design '92, Kluwer Academic Publishers.