An Agent Design for Effective Negotiation Dialogues

Bryan McEleney and **Gregory O'Hare**
Department of Computer Science
University College Dublin
Dublin 4,
Ireland
{bryan.mceleney, gregory.ohare}@ucd.ie

Abstract

A design is presented for a negotiating agent that can construct coherent joint plans with human or artificial agents. In negotiation there is always a trade-off between plan quality and dialogue length. In dynamic conditions and with human partners, length becomes critical. The approach to efficient negotiation is to use an acquaintance model that predicts which plans will be acceptable. The negotiation dialogue then consists of exchanges to construct the acquaintance model and exchanges of plan proposals.

1. Introduction

Planning problems involve reasoning about the effects of actions so that future actions can be chosen to satisfy a set of input goals. In a multi-agent system, planning is a process that occurs simultaneously in a number of agents, each selecting their own local plans, and for the most part independently. However, the agents are situated in the same environment and interactions between individual plans occur which must be identified and reasoned about. Such interactions might be destructive, where one agent's chosen action has an effect that negates another's goal, or they might be constructive, where one agent's action can satisfy the other's goal. In each case the agent must identify and reason about the plans of other agents in the community as well as about its own. Better still, agents can exert some influence on one another's choices by participating in a negotiation dialogue.

This paper considers the design of a negotiating dialogue agent to conduct plan negotiations in practical situations with other artificial agents or with a human user through a natural language interface. The multitude of plan options, the irregularity in plan quality over the search space, and the uncertainty in which plans are agreeable to other agents can make the search for an acceptable joint plan a lengthy one. Dialogue length is of critical importance for two reasons. First, the dialogue partner might be a human user with bounded reasoning resources in terms of evaluating and discussing proposals. Second, planning must often be done in the real world, with continually changing goals and circumstances that must be responded to in a timely manner. The aim is thus to construct a negotiating agent that can reach an acceptable proposal within as few exchanges as possible, and with little computational burden upon itself and other agents in generating and evaluating proposals. It is proposed that the negotiating agent should incrementally construct an acquaintance model, capturing the preferences of the other agent by constructing the topology of the plan search space with respect to its utility. Exchanges in the negotiation are then of two main types, one that maps out the plan search space by explicitly discussing goals and preferences, and one that involves proposal and evaluation of plan instances.

Examples from human-human dialogue corpora (Allen et al 1995, Kowto and Price

1992) demonstrate that acquaintance models are used to some extent to make negotiation dialogues as brief as possible. The types of utterance that occur in the dialogue can be analysed and shown to be part of either acquaintance model construction, or actual negotiation exchanges. The corpus data sets out the rules of play in the dialogue, and is a necessary input to the design of a dialogue system that can communicate fluently and flexibly in natural language, and in a manner similar to human negotiators. The system offers a model of practical negotiation in real-time environments that is inspired by human performance in planning problems and by ideas about multi-agent coordination for artificial agents.

2. Levels of cooperation

Acquaintance models represent the mental state of other agents, in terms of mental attitudes such as beliefs and desires, and plan related states such as actions already chosen and committed to, capabilities of executing actions and available resources for future commitments. They can be constructed in part from observation of behaviour, in part from explicit dialogue, and in part from plan inference drawn on these two, which hypothesises mental states consistent with underlying spoken and domain actions. There are some situations in which an acquaintance model can be formed purely from observation of domain actions with no use of communicative acts. For example, when two agents individually construct a plan involving the use of the same corridor, the potential interaction of each blocking the other's way can be identified through a process of observation and plan recognition as the plan is executed. Plan recognition is a useful strategy in situations where communication ability is limited or in hostile *zero-sum* situations, such as interactions between opponents in a robot soccer game where one agent's goal is the negation of the other's. In the hostile situation, agents hide their plans and preferences from their opponent since being able to predict the

other's actions produces more coherent action, and in the zero-sum situation this implies a worse utility for the agent. In cooperative situations the utility gain of one agent's action towards another's goal can be greater than the utility cost to itself, motivating profitable plan agreements between agents. In this situation, agents will discuss their preferences with others so that profitable interactions can be found, resulting in each agent adopting a plan that has a greater utility than the plan they would have constructed without negotiating. It is interesting to note that even in cooperative situations, speakers will rely on the hearer's ability to recognise their preferences through plan recognition, by providing just enough evidence in their utterance to disambiguate their meaning (Allen, Perrault, 1980) (Grice, 1975). In this paper, negotiations between self-interested agents are discussed − a type of agent that is cooperative with others as long as it can come to deals that produce a mutual increase in utility. Such agents will freely discuss their mental state so that joint plan opportunities may be identified and plans constructed. Self - interested agents represent common negotiation situations such as a community of agents that trade goods and services with one another.

3. The negotiation process

The negotiation process is outlined as follows. The negotiating agent begins the planning process with a space of joint plans, S, from which to choose a plan, and a utility function that ranks those plans according to the degree to which each satisfies the agent's goals. A joint plan is the combination of the agent's own plan and that of its acquaintances, that is, a set of actions each associated with the agent that is to execute them, and a set of ordering constraints. The object of the negotiation process is to select the highest-ranked joint plan possible, but within constraints of acceptability imposed by the acquaintances. At the start of the process, little is known about the constraints, but through explicit sharing of information and as proposals are exchanged and evaluated, the agent can construct an

accurate map of the acceptability of its plan space. Each agent starts the negotiation by exploring regions in the plan space of high utility, and then bargains its way towards regions that are of lower utility, but are more likely to be acceptable to the other agents. Once an acceptable plan is found, the agents enter a subdialogue in which they commit to the shared plan. The commitment is a trusted agreement whereby each agent agrees to carry out their portion of the plan to the best of their ability. At this point each of the agents may execute their portion of the complete plan.

4. Complexity and resource bounds in negotiation

Constructing plans can be a time-consuming process due to the combinatorial complexity of the plan search space. In some cases, it is an intractable problem, and the combinatorial increase in generating a multi-agent plan can aggravate this situation. The result of the planning process can be a correspondingly large selection of plan proposals, of uncertain utility, which are the objects in the negotiation. It might be argued that the agent should reduce the complexity by decomposing the problem and negotiating only about parts of the plan at time. This is the approach taken with the contract net (Smith, Davis, 1983). In general though, actions within plans interact by satisfying one other's preconditions and by consuming the same resources. These interactions prevent the negotiating agent from dividing the plan into a set of n independent subplans that can be individually negotiated. With such a large space of plan proposals to consider, a blind search for agreeable plans may be necessary to guarantee optimality, but could be very expensive in terms of the time required for each proposal to be spoken, understood and evaluated, especially for a human user with bounded reasoning resources. Instead the negotiating agent should take great care in predicting which proposals yield a high utility for the other agent.

Another difficulty that occurs in multi-agent planning is that the actions may be too complex to model accurately, or may be uncertain in their effects. Unpredictable events can happen in the environment, including the actions of other agents who plan their actions dynamically. One agent's model of another is inherently incomplete and agents who appear at one moment to be agreeable towards a plan may change their desires unexpectedly. For example, consider where one agent plans to use a resource belonging to another agent. If another agent becomes interested in the same resource, the owner's utility value in leasing the resource decreases to correspond with increased demand. The negotiation model must accommodate such uncertainties by interleaving negotiation with plan execution. Agents must participate in a continuous dialogue where opportunities continually arise as circumstances change, and commitments to plans must be revoked and renegotiated when it becomes apparent that plans are no longer profitable. An example of such a dynamic situation is a robot soccer game. Players may negotiate general strategies which can be fixed before the game begins, but must also adapt and make short-term tactical decisions as the game progresses. In dynamic situations agents must keep their negotiations as short as possible so that they can act before their plan becomes inapplicable.

5. Constructing and using an acquaintance model

To produce acceptable joint plans within a reasonably short dialogue, an acquaintance model is constructed during the dialogue that can be used to estimate the utility value of plan proposals to the other agents. Candidate proposals can then be found without engaging in the expensive communicative process of making a proposal and asking the other agent to evaluate it. Instead, the acquaintance model is used as a utility approximating function that can filter those proposals that are more likely to be accepted by the acquaintance.

The acquaintance model is a tuple

<C, B, D, P, PS, U: PS → R> where:

C is the agent's capabilities, a combination of the set of actions the acquaintance is capable of executing, the resource requirements of those actions, and the amount of each resource available to the agent

B is the set of acquaintance beliefs

D is the set of acquaintance desires

P is the partial plan (if any) that the acquaintance has already committed to executing. Further negotiation must accommodate P.

PS is the plan space, the set of joint plans that can satisfy D. Each joint plan is a partially ordered set of actions to be executed by the agent and the acquaintance.

U is the utility estimate, a function that maps plans onto the real numbers, which ranks the plans within the plan space according to the acquaintance's utility

In addition there is a function F: PS → R which evaluates a plan with respect to the agent's own utility

Goal description phase

Using an acquaintance model, the dialogue is composed of two major phases of interaction. The first, the called *goal description phase* is where the agent and the acquaintances try to characterise one another's utility function, capabilities, plans, desires and beliefs. Such subdialogues can be initiated either by the agent or the acquaintance, and generally occur at the start of the dialogue. Two subtypes of such interactions can be identified. First, the agents exchange desire and belief descriptions, which describe the world states the agent would like to achieve, and the perceived state of the world at the current point

in time. Desire and belief descriptions define the space of plans PS. For example, if one agent states "I need to make some travel arrangements", PS can be defined as a space of plans to buy travel tickets. The acquaintance may also have constructed and committed to executing certain actions, P, which must be included in all of the plans in PS. Where these are commitments made with another agent, revoking might be an especially expensive process, and these actions are considered non-negotiable. For example, the agent may state "I need to book a train that will connect with the flight I booked with agent J".

Within the space of potential plans, some plans are considered better than others by the acquaintance, and so the second interaction subtype serves to construct the U, the utility estimate function. These interactions can describe hard constraints such as "I must travel first class", or more fuzzy ones such as "I prefer to fly early in the morning". Even naive dialogue system users are familiar with the goal description process, and know how to provide a description that is just succinct and detailed enough that the system should be able to provide a suitable proposal. In the SRI transcripts (Kowto, Price 1992), for example, the user will often take the initiative in the dialogue by specifying the constraints that most distinguish the kind of plan they would be interested in, such as which airports they wish to fly between and which dates they wish to fly on. Often though the user will neglect to mention some features, such as for example whether an aisle or window seat is required, either through forgetfulness or by assuming such a feature could not be adjusted within the plan. The system must accommodate such behaviour by taking the initiative and asking the user to characterise his utility with respect to the remaining plan features. This is especially important where the feature values have a significant effect on the system's plan utility.

An important phenomenon occurs in the goal description phase, which is a result of the acquaintance trying to be as brief as

possible. This is that the acquaintance tends not to explicitly state desires, plans and preferences. Instead, the acquaintance expects the agent to infer and adopt the *most likely* set of mental attitudes that are consistent with his utterance, unless otherwise stated. For example, by stating "I want to book a holiday in Spain", the speaker expects the hearer to recognise the most common plans and desires, such as the speaker travelling to Spain, visiting the beach, and staying for a week or two. To make these inferences, a domain-specific plan library can be used, associating desires with plans, to indicate the most likely plans to realise a desire, and the most likely desires given a plan or plan fragment.

Negotiation phase

At some point in the dialogue, one of the agents deems that the goal description process has gathered enough information to begin proposing specific plans that have a reasonable likelihood of being accepted by the other. This is where the *negotiation phase* of the dialogue begins. To determine when this point is reached, the agent uses its knowledge of how detailed its acquaintance model has become in the goal description process, and its knowledge of how detailed the other agent's acquaintance model has become. The negotiation phase begins with the proposal of a particular plan which is communicated to the acquaintances. To come up with a plan, the agent invokes its planner on the combination of its own goals and the acquaintance model. The search is guided by the agent's own utility function, F, and by the utility estimate for the acquaintance, U. The plan is chosen from among those that yield a positive utility gain for both the agent and the acquaintance. The acquaintance then evaluates the plan with respect to its utility function. At this point the acquaintance may accept the plan at which point the dialogue finishes. Alternatively, the plan may be rejected. The acquaintance may explain why the plan was not accepted, allowing an update to U and a further iteration of the negotiation phase, or it may declare that

its utility gain is not sufficient to warrant an agreement, or it may update its own acquaintance model with the evidence that the agent has a preference towards its proposal, and then produce a counterproposal plan. The negotiation phase can be thought of as a generate-and-test hill-climbing search, where the evaluation of each proposal contributes to the utility estimate U, and each agent searches for proposals that have a high value for F and U and are similar to proposals presented earlier in the negotiation.

Commitment and revoking phases

During the agent's lifecycle, negotiations may occur a number of times. With new opportunities and unexpected changes in the environment, agents continuously make new commitments and may revoke existing ones. Part of the dialogue structure represents where an agent identifies that a commitment is no longer appropriate, and attempts to revoke it. An abstract dialogue grammar for negotiating agents is as follows:

LifeCycle = Cycle*

Cycle = NegotiationCycle

Cycle = NegotiationCycle Revoke

NegotiationCycle = NC* Commit

NC = Goal_description_phase

NC = Negotiation_phase

5. An Example

An excerpt from the SRI transcripts (Kowto, Price 1992) illustrates the goal description and negotiation processes in a human-human dialogue. A fragment of the dialogue can be paraphrased as:

A: I believe there's an eight o'clock United flight on May 12th?

B: The fare is one hundred and ninety-eight dollars

A: What happened to the seventy-eight dollar fare?

B: For those fares you need to stay over a Saturday night

At the start of the dialogue A has a plan space and approximate utility function for B based on an earlier plan proposal, and reasons that a similar proposal will have a similar utility with respect to B. He therefore deems that entering a negotiation phase is appropriate, and suggests this new proposal expecting that it will be acceptable to B. However the proposal is rejected when B's utility is found to be less than expected. A goal description subdialogue updates A's acquaintance model to account for plans that involve a Saturday night stopover. At this point, the negotiation could have continued if there were alternative proposals, but no alternative is presented by B, so A accepts and a deal is reached.

6. System Design

The design of the negotiation system is based on the family of BDI agent architectures, of which the Grate (Jennings, 1993) and IRMA (Bratman et al, 1988) are examples. These architectures have been designed to work in dynamic environments, under conditions of bounded rationality, where plans must be quickly and continually updated in response to environmental changes. BDI architectures are centred around a combination of a belief set, B, which is what the agent holds to be true about its environment, a set of desires D, which are the goals or preferences of the agent, and a set of intentions, I, which is the goals and plans that the agent has committed to executing. The agent executes a control loop where first the beliefs are updated from sensor data. Then the agent's intentions are updated by selecting a consistent subset of the available desires, and a plan is formed to realise those intentions. Finally, actions that must be executed at the current cycle are carried out, and the agent returns to the top of the loop. Where possible, and particularly where plans interact with those of other agents, the BDI agent maintains a level of commitment towards its selected plan. Figure 1 outlines the architecture of the system.
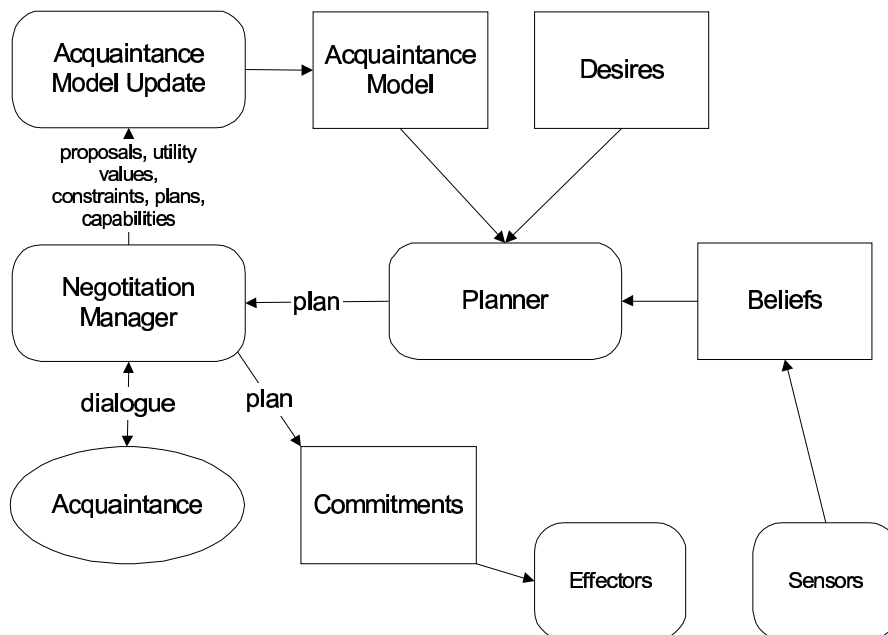

Figure 1. Negotiation Agent Architecture

The planning component of the traditional BDI architecture is coupled with the negotiation manager, which is responsible for controlling the negotiation dialogue. During the goal description phase, the negotiation manager updates the acquaintance model with the beliefs, capabilities, desires and partial plan expressed by the acquaintance. The utility function, U, is implemented as a combination of heuristic constraint rules derived from the goal description phase of the dialogue, and specific evaluated proposals expressed in the negotiation phase. These are used to form a case-base of high quality plans from which new proposals can be adapted. The negotiation manager repeats a negotiation loop, where it calls the planner to produce a new proposal, the proposal is communicated to the acquaintance and a response received. The planner takes as its input both the acquaintance model and its own desires, and produces a plan that satisfies both. As the negotiation proceeds, the planner balances its utility in favour of the acquaintance, so that an acceptable plan is reached quickly. Once the negotiation is finished, the agent's commitment set is updated and commitments scheduled for the current time-point are executed.

7. Evaluation and Future Work

In evaluating a plan negotiation system, the two factors of plan quality and dialogue length must be measured. There is a trade-off between these, and so a weighted sum of the two provides one suitable quantitative metric. In addition, the system should perform well with naive users, communicating fluently in natural language rather than presenting and interpreting the dialogue in an artificial language. So far, the natural language capability has been put aside so that the core problem of controlling the dialogue can be developed. Indeed it is possible to implement and evaluate dialogue strategies without using a natural language interface, so the short-term plan is to do just that, while in the long term, the issues of generating and understanding the

types of utterances that occur in negotiation dialogues will be dealt with. With a complete dialogue control system and natural language interface, human-human dialogues can be compared with human-system dialogues in the same experimental setting.

With the objective of producing a system that behaves like a human-being would, and bearing in mind that human-beings can be accomplished negotiators, further corpus analysis will help to highlight how dialogue control should be approached. Dialogue structure can be analysed to identify the phases of negotiation and their content. Corpora already exist in for example the SRI transcripts in the ATIS project (Kowto, Price 1992), which concern dialogues with a travel agent, the TRAINS project (Allen, 1995), which concern dialogues between a human manager and a planning assistant in a planning task, and in the Verbmobil corpus, which concerns negotiations about meeting scheduling. While such corpora are useful, other experiments may be carried out to investigate how humans cope with negotiation and planning, particularly under pressing time constraints. A suitable task might be a simulated soccer game, a blocks-world task, the Tileworld (Pollack, Ringuette 1990), or the postman problem (Zlotkin, Rosenschein, 1989)

The idea of that acquaintances can be grouped into stereotypes classes with similar properties is important in a negotiation system, as a means of reducing the length of the goal description phase. For example, by contacting a travel agent, an agent can bring to mind stereotypical properties constructed from past experience with travel agents, such as their capabilities, and typical utility values associated with their plans. Such information can be used to initially construct the user model, which can then be further refined in the dialogue. On the travel agent's side, asking whether the customer is a businessman or a student for example can influence the acquaintance model. An extra dialogue phase of stereotype activation might

be added to the dialogue structure, and evaluated to see if it can produce better plans in the same dialogue length.

References

James Allen, C. Raymond Perrault. 1980. *Analysing Intention in Utterances*. Artificial Intelligence, 15(3): 143 - 178

James Allen et al. 1995. *The TRAINS project: A Case Study in Building a Conversational Planning Agent*, Journal of Experimental and Theoretical Artificial Intelligence 7: 7 – 48.

Michael Bratman, David Israel, Martha Pollack. 1988. *Plans and resource-bounded practical reasoning.* Computational Intelligence 4: 349 – 355.

Edmund Durfee and Victor Lesser. 1991. *Partial Global Planning: A Coordination Framework for Distributed Hypothesis Formation.* IEEE Transactions on Systems, Man, and Cybernetics 21(5) : 1167-1183.

Richard Fikes and Nils Nilsson. 1971. *STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving.* Artificial Intelligence 2: 189 – 208.

Herbert Paul Grice. 1975. *Logic and Conversation.* In Syntax and Semantics 3: Speech Acts Cole and Morgan, eds :41 – 58. Academic Press, New York, NY.

Nick Jennings. 1993. *Specification and Implementation of a belief-desrire-joint-intention Architecture for Collaborative Problem Solving.* Int J. Intell. Coop. Inf. Syst., 2 (3): 289 - 318 .

Martha Pollack, Marc Ringuette. 1990. Introducing the Tileworld: Experimentally Evaluating Agent Architectures. Proceedings of the Eighth National Conference on Artificial Intelligence. Boston, MA, USA, August 1990

Tuomas Sandholm, Victor Lesser. 1995. *Issues in Automated Negotiation and Electronic Commerce: Extending the Contract Net Framework.* In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS'95) : 328-335.

Reid Smith and Randall Davis 1983. *Negotiation as a metaphor for distributed problem solving.* Artificial Intelligence 20 (1) : 63 –109.

Jaqueline Kowto and Patti Price. 1992. *SRI Transcripts.* Transcripts derived from audio-tape conversations made at SRI International, Menlo Park, CA.

Gilad Zlotkin and Jeffrey Rosenschein. 1989. *Negotiation and Task Sharing Among Autonomous Agents in Cooperative Domains.* Proc. 11[th] IJCAI, Detroit, Michigan, USA, August 1989. 912- 917