

Architectures for negotiating agents

Ronald Ashri¹, Iyad Rahwan², and Michael Luck¹

¹ Dept of Electronics and Computer Science, Southampton University,
Southampton, UK

ra00r@ecs.soton.ac.uk

² Department of Information Systems, University of Melbourne
Melbourne, Australia

i.rahwan@pgrad.unimelb.edu.au

Abstract. Automated negotiation is gaining interest, but issues relating to the *construction* of negotiating agent architectures have not been addressed sufficiently. Towards this end, we present a novel agent construction model that enables the development of a range of agent architectures based on a common set of building blocks. In this paper we identify the fundamental components needed for two generic classes of negotiating agents: *simple negotiators* and *argumentative negotiators*, and use our model to describe them. We demonstrate how the model allows us to reuse fundamental components across these negotiation architectures.

1 Introduction

In multi-agent environments, agents often need to interact in order to achieve their objectives or improve their performance. One type of interaction that is gaining increasing interest is *negotiation*. We adopt the following definition of negotiation that reconciles views proposed by [6] and [13], which we believe is a reasonable generalisation of both the explicit and implicit definitions in the literature.

Negotiation is a form of interaction in which a group of agents, with conflicting interests and a desire to cooperate, try to come to a mutually acceptable agreement on the division of scarce resources.

Agents typically have conflicting interests when they have competing claims on scarce resources, which means their claims cannot be simultaneously satisfied. Resources here are taken to be very general. They can be commodities, services, time, etc. which are needed to achieve something.

To address this problem, a number of interaction and decision mechanisms have been presented³. There has been extensive work on implementing frameworks of negotiation based on auction mechanisms (as evident, for example, in the Trading Agent Competition [12]) and frameworks that adopt heuristic-based bilateral offer exchange

³ For a more comprehensive comparison between different approaches to negotiation, the reader can refer to [6].

(e.g. [4, 5]). Recently, argumentation-based approaches [7, 8, 10] have been gaining interest. However, there are very few implemented systems that cater for more sophisticated forms of interaction such as argumentation. One of the reasons for this is that many of these frameworks involve complex systems of reasoning based on logical theories of argumentation, for which there are still many open research questions [9]. Another reason is that there are no software engineering methodologies that structure the process of designing and implementing such systems. This is why in most cases, these systems are implemented in an *ad hoc* fashion.

The aim of this paper is to address the software engineering issues related to the development of architectures for negotiating agents, ranging from simple *classical* agents to more complex *argumentative* negotiators. More specifically, this paper advances the state of the art in automated negotiation in the following ways. First, it presents a novel agent construction model that enables the description of a range of agent architectures through a common set of concepts and building blocks. Secondly, it uses this agent construction model in conjunction with a general negotiation framework to design and describe the architectures of two generic classes of negotiating agents: *simple negotiators* and *argumentative negotiators*. The paper demonstrates how a generic architecture for argumentative negotiators can be achieved by extending the simple negotiator architecture and reusing its components, and shows how this modularity is facilitated by the construction model.

We begin by presenting the agent construction model in Section 2, supported by a brief analysis of the engineering requirements the model seeks to fulfill. In Section 3, we discuss the negotiation framework before presenting a generic architecture for a basic negotiating agent. We then explain how the construction model allow us to *re-use* this basic negotiation architecture in developing a more complex architecture for an agent performing argumentation-based negotiation.

2 Engineering Agent Architectures

In this section, we present the design approach that we have applied to the specification of architectures for negotiating agents. In essence, we require an agent construction model that allows the description and development of a range of architectures. In order to achieve this we have two possible avenues to explore. One option is to define a generic agent architecture and describe other architectures in terms of this generic architecture. The drawback of this approach is that there may be features of other architectures that cannot be directly *translated* to the generic one. The second option, which we follow, is to provide an architecturally neutral model, so as to avoid this translation problem. The challenge is to provide a model that is specific enough so that it actually aids in the construction of agents. Through such a model we can view a range of architectures based on a common view of agents without any loss in expressive capability. Furthermore, any resulting model must also allow for the modular construction of agents in order to meet general software engineering concerns and to delineate clearly the different aspects of an architecture, as we discuss below. Such a fine-grained approach leads to a better understanding of the overall functioning of the agent as well as how it can be altered. Finally, we need to be able to re-configure the resulting ar-

chitectures easily if possible, even at run-time, in order to deal with dynamic, complex dependencies that develop in heterogeneous computing environments.

SMART The agent construction model is based on SMART [3] (Structured, Modular Agent Relationships and Types), which provides us with the foundational agent concepts that allow us to reason about different types of agents, and the relationships between them, through a single point of view. We chose SMART because it provides us with appropriate agent concepts without restricting us to a specific agent architecture. Furthermore, SMART has already been successfully used to describe several existing agent architectures and systems (e.g. [2]).

We avoid a more complete presentation of SMART and focus on just those concepts that are used for the agent construction model. In essence, SMART provides a compositional approach to the description of agents that is based on two primitive concepts, *attributes* and *actions*. Attributes refer to describable features of the environment, while actions can change the environment by adding or removing attributes. Now, an agent is described by a set of *attributes* and a set of *capabilities*, where capabilities are actions an agent can perform. An agent has *goals*, where goals are sets of attributes that represent desirable states of the environment for the agent. On top of this basic concept of an agent, SMART adds the concept of an *autonomous agent* as an agent that generates its own goals through *motivations*, which drive the generation of goals. Motivations can be preferences, desires, etc., of an autonomous agent that cause it to generate goals and execute actions in an attempt to achieve those goals.

This approach to agent description fits well with our requirement for architecture neutrality but does not sufficiently address our requirements for modularity and run-time reconfiguration. We address these issues via a decoupling of the different aspects of an agent and a component-based approach to agent construction. Both these aspects of the construction model are described below.

Decoupling description, structure and behaviour In this subsection, we describe how we extend SMART to provide a more flexible decoupling of agent aspects. SMART allows systems to be specified based on an observer's point of view, in terms of their attributes and goals, as well as the actions they can perform. However, this description does not show how agents are built or how they behave. In other words, the focus is on the *what* and not the *why* or *how*. We call this a *descriptive specification*, since it essentially describes the agent without analysing the underlying structures that sustain this description. Along with the descriptive specification we need to have the ability to specify systems based on their structure, i.e. the individual components that make up agents, as well as their behaviour. Thus we extend SMART by the addition of a *structural specification* and a *behavioural specification*.

The structural specification enables the identification of relevant building blocks and how different sets of building blocks enable the instantiation of different agent types. The behavioural specification of an agent addresses the process through which the agent arrives at such decisions as what actions to perform. These views, along with the descriptive specification, can provide a more complete picture of the system. The agent

construction model, described next, reflects these concepts by allowing direct access to these different aspects of agents, based on a clear decoupling at the architectural level.

2.1 Agent construction model




Fig. 1. Agent shell

The basic principles of the model are illustrated in Figure 1. A *shell* acts as the container in which *components* are placed. It manages the sequence in which components execute and the flow of information between components. Control policies relating to the permissions an agent has in a specific environment are defined within the shell in order to make them independent of the agent architecture. Finally, attributes describing the agent as a whole are defined as part of the shell.

Components encapsulate specific types of actions that an agent can perform and are grouped into four categories. *Sensors* (dashed rectangles) receive information from the environment, *infostores* (rounded corner rectangles) store information, *actuators* (continuous line rectangles) perform actions that affect the environment and *controllers* (accented rounded corner rectangles) are the main decision-making components. Controllers analyse information and delegate actions to other components. By dividing components into these categories we can abstract between high level design, providing an understanding of an architecture early, and low level design, where specific mechanisms for controllers, sensors or actuators need to be defined. Each component is described using two types of attributes: *stateless attributes* refer to persistent characteristics, such as the kind of communication language the agent uses, while *situation attributes* refer to attributes describing the component's current state (e.g. the parties with whom the agent is currently negotiating).

Information flows through *links* that the shell establishes between components. Links are uni-directional, one-to-one relationships. The information that flows through links between components is packaged within *statements*. One component acts as the producer of a statement and the other as the consumer. Statements are typed, and although currently just two types are defined, *INFORM* and *EXECUTE*, designers may choose to define different ones depending on application needs. Inform-type statements are used when one component simply notifies another component about something, while execute-type statements are used when a component wants an action to be performed by another component. All statements are divided into a *body* and *predicates*.

The body carries the main information (e.g. an update from a sensor), while the predicates carry additional information (e.g. the source of information or specific conditions associated with the execution of the action).

The sequence in which components execute is defined as the *execution sequence* of the architecture. Execution of a component includes the processing of statements received, the dispatch of statements and the performance of any other needed actions. The execution sequence is an essential aspect of most agent architectures and, by placing the responsibility of managing the sequence within the shell, we can easily reconfigure it at any point during the agent's operation.

Agent design begins with an empty shell. It is then specialised by defining control policies in order for it to meet application requirements or the demands of the environment within which it will operate. Then, shell-specific attributes can be defined to form part of the description of the agent to the outside world. The components can then be loaded into the shell, and links, as well as an execution sequence, can be defined. With the execution sequence in place, the operational cycle of the agent can begin. The agent lifecycle can be suspended or stopped by stopping the execution sequence and can be modified by altering the execution sequence, modifying relationships between components, or by applying alternative control policies.

One of the main benefits of this approach is that it is possible at any moment to have access to the three views of agents as described previously. The *descriptive specification* can be obtained by aggregating the situation attributes and stateless attributes from each component as well as the attributes contained in the shell; the *structural specification* is given by the components themselves; and the *behavioural specification* is given by the execution sequence and the links between components.

3 Negotiating Agent Architectures

With the agent construction model in place we can investigate the suitability of our model for specifying flexible negotiating agent architectures. Before we start describing negotiating agents, however, we discuss the main components of a *negotiation framework*. In addition to the negotiating agents, a negotiation framework usually includes a communication language and an interaction protocol. For example, a negotiation framework based on a simple English Auction protocol would need a communication language locution (or performative), say *propose(.)*, that can express bids. The protocol is the set of rules that specify, at each stage of the interaction, what locutions can be made, and by whom. In addition, the framework needs a language for representing information about the world, such as agents, agreements, arguments, and so on. This information is used within the communication language locutions to form utterances. For example, a bid might be presented as *propose(a, b, {toyota, \$10K})*, where *a* and *b* are the sending and receiving agents, and *{toyota, \$10K}* is the specification of the proposal. Finally, a negotiation framework usually includes various information stores needed to keep track of various information during the interaction. This information may include proposals made by different agents, concessions they have committed to [13], and so on. Finally, the framework also needs a set of additional non-protocol rules.

These may include rules that identify the winner in a particular negotiation, or rules that specify that agents cannot retract their previous proposals, and so on.

In this paper, we focus our attention on the construction of the agents within the framework. So we do not address, for example, how protocols can be specified in a modular fashion (this has been investigated in [1] for example), or how the locutions can be verified. We assume that developers have at their disposal definitions of the appropriate negotiating protocols, domain ontologies and communication languages, and instead deal with the problem of framing such mechanisms within an appropriate agent architecture. Note that we do not claim to have specified the *only* way of describing negotiating agents. Instead, we attempt to provide *a* construction model that is generic enough to capture a variety of negotiators.

3.1 Basic Negotiating Agent

Basic negotiating agents include those participating in auctions or those engaged in bilateral offer exchanges. The common aspect of these agents is that they engage in interactions in which the primary type of information exchanged between agents are proposals (i.e., potential agreements). We call the agents *basic* in order to distinguish them from agents that can engage in more sophisticated forms of negotiation which allow the exchange of meta-information (or arguments). The proposed architecture for basic negotiation agents is illustrated in Figure 2 and is described through the three different views below. We follow the conventions described earlier for illustrating the different types of components and the connecting arrows illustrate the flow of statements.

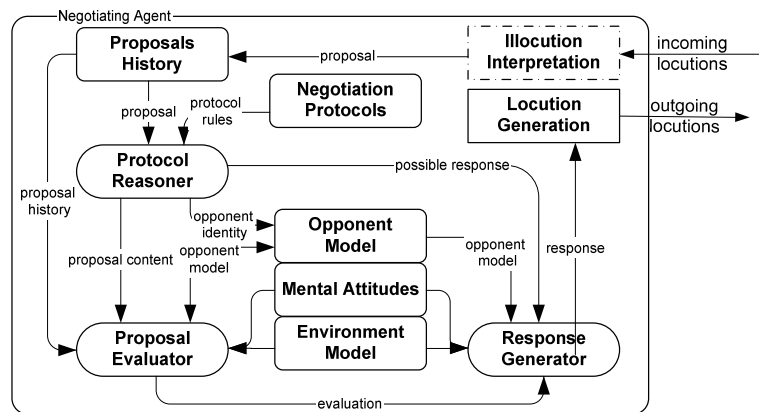


Fig. 2. Negotiating Agent Architecture

Descriptive Specification The description of the negotiating agent is based on its attributes, capabilities, goals and motivations. The goals of the agent, i.e. the desired negotiation outcomes, can be represented in the *Mental Attitudes* infostore, and these refer to specific application domains. However, the architecture does not require explicit representation of agent goals. We could have as an overarching goal the achievement of the

environmental state that represents the desired negotiation outcome for the agent. This desired state would be determined by the mechanisms used by the *Proposal Evaluator* and *Response Generator* components, which ultimately decide when this environmental state has been reached. Here we see how access to an overarching, architecturally neutral, agent model allows us to reason about such things as goals even though they find no explicit representation in the architecture. Attributes of the agent are given by the information that is stored and interpreted inside components, and flows between them through statements. These attributes include representations of beliefs about the opponents, the environment, mental attitudes, negotiation protocols, and so on. Because all this information is explicitly represented within components, and stateless and state-dependent information is separated, we can easily extract it. The capabilities of the agent are given by the aggregation of capabilities of each component and can be understood, in our case, by referring to individual components in the architecture diagram. This is possible because our architecture attempts to represent the main capabilities with separate components so as to make clear the various functionalities required. However, alternative designs could (as in many implementations in the literature) combine a number of components (e.g. the representation of opponents, mental attitudes and environment) within a single component. In such cases, the descriptive specification remains unaltered, since the capabilities exist, but the structural specification refers to different components that combine those capabilities. Finally, the motivations of the agent, if the agent were autonomous, form part of the agent's mental attitudes, and ultimately guide the agent's decisions. Exactly how these are defined depends on the application. In many auction-based mechanisms, for example, the motivations are represented in the form of a utility function.

Structural specification The structure of the agent refers to the components that make up the architecture. Messages are received, checked and parsed through the *Illocution Interpreter*. The *Proposals History* infostore keeps track of the various proposals received, while the *Negotiation Protocols* infostore contains the rules relating to the negotiation protocols. By separating the rules dictating the protocol from the reasoning about the protocol we can more easily extend the agent to handle different protocols. The *Opponent Model* infostore keeps track of opponent models, while the *Environment Model* maintains information about the environment in which the agent is situated. Information such as the agent's preferences is stored in the *Mental Attitudes* infostore. Decisions are taken by three controllers that, abstractly, support the different negotiation stages. The *Protocol Reasoner* checks whether the proposal received by the opponent is a valid response based on the negotiation protocol. The *Proposal Evaluator* evaluates the proposal and the *Response Generator* generates an appropriate response based on this evaluation. Finally, the *Locution Generator* packages responses in the appropriate message format and handles outgoing communication.

Behavioural Specification The behaviour of the agent is largely dictated by the flow of information through the architecture. It begins by message interpretation and storage in the *Proposals History*. The current proposal and information of the history of proposals is sent to the *Protocol Reasoner*, which uses rules in the *Negotiation Protocols* infostore to check the validity of the proposal. If it is valid it is forwarded to the *Proposal Evaluator*, which retrieves information about the opponent from the *Opponent Model*

infostore. This controller uses this information along with information from the *Mental Attitudes*, *Environment Model* and *Proposal History* to evaluate the proposal. The result of the evaluation is sent to the *Response Generator*, and the opponent model may be updated. This controller also uses information from the now updated opponent model, the mental attitudes and environment model in order to generate a response. It also takes into account the negotiation protocol rules in order to generate the appropriate response. The response is packaged in the appropriate format by the *Locution Generator* before it is sent to the opponent.

3.2 Argumentative Negotiating Agent

Here, we instantiate the architecture of the basic negotiating agent in order to provide a generic description of agents capable of conducting argumentation-based negotiation (ABN). An argumentative negotiator shares many components with the basic negotiator. For example, it also needs to be able to evaluate proposals, generate proposals and so on. What makes argumentative agents different is that they can exchange meta-information (or arguments) in addition to the simple proposal, acceptance, and rejection utterances. These arguments can potentially allow agent to (i) justify their negotiation stance; or (ii) influence the counterparty's negotiation stance [7]. This may lead to a better chance of reaching agreement; and/or higher-quality agreements. In ABN, influencing the counterparty's negotiation stance takes place as a result of providing it with new information, which may influence its mental attitudes (e.g., its beliefs, desires, intentions, goals, and so on). This might entice (or force) the agent to accept a particular proposal, or concede on a difficult issue. Arguments can range from threats and promises (e.g. [11]) to logical discussion of the agent's beliefs (e.g. [8]) or underlying interests [10].

In order to facilitate ABN, the logical and communication language usually needs to be capable of expressing a wider range of concepts. For example, the proposal might instead be represented as *propose*(*a*, *b*, *P*, *A*) where *a* and *b* are agents, *P* is a proposal, and *A* is a supporting argument denoting why the recipient should accept that proposal. ABN frameworks may also allow agents to explicitly request information from one another. This may be done, for example, by posing direct questions about agent's preferences or beliefs, or by challenging certain assumptions the agent adopts. Since in this paper we are more interested in the abstract structures within the agents, we shall not discuss these issues in more detail. In order to be capable of engaging in ABN, an agent needs the following additional capabilities:

1. **Argument Evaluation:** This component encompasses the ability of the agent to assess an argument presented by another, which may cause updates to its mental state. This is the fundamental component that allows negotiators' positions to change.
2. **Argument Generation:** This component allows the agent to generate possible arguments, either to support a proposal, or as an individual piece of meta-information. The locution generated may also be a question to present to the opponent.
3. **Argument Selection:** Sometimes, there might be a number of possible arguments to present. For example, an agent might be able either to make a promise or a threat to its opponent. A separate component is needed to allow the agent to choose the more preferred argument. Selection might be based on some analysis of the expected influence of the argument, or on the commitments it ties the utterer to.




Fig. 3. Argumentation-based negotiation agent Architecture

Figure 3 shows the specification of an argumentative agent using our construction model. All components from the basic negotiating agent have been used, complemented by the additional capabilities needed for ABN. Note that the diagram has been simplified for clarity (e.g. a bidirectional link stands for a pair of unidirectional links). Furthermore, the link from *Negotiation Protocol* to *Response Generator* and *Argument Generator* has been omitted although it is, of course, necessary. Below we show how the descriptive and behavioural specification are changed. The structural specification changes by adding the three new components that deal with ABN.

Descriptive Specification A crucial difference between the simple negotiation agent and the ABN agent is that arguments from opponents can change the agent's mental attitudes so the agent's goals or motivations may change based on the new information obtained. As a result, even this aspect of the descriptive specification is dynamic and the ability to refer to this changing descriptive specification directly, at run-time, by extracting the relevant attributes is crucial. The descriptive specification must also include the new decision-making capabilities of the agent.

Behavioural Specification Here the flexibility provided by the agent construction model is particularly evident. The agent essentially has the same links and information flows, and is simply *extended* with links to the new controllers and refined through changes to the execution sequence. The opponent model, mental attitudes and environment model are now updated by the evaluation of the argument received before the proposal is evaluated. The response is not sent directly to the opponent but arguments may be attached to the proposal by the *Argument Generator* and *Argument Selector* components. Finally, both the *Response Generator* and *Argument Generator* use the negotiation rules in order to determine what type of responses are possible.

4 Conclusions and Future Work

Automated negotiation is gaining increasing interest, but this growth is not matched by appropriate agent-oriented software engineering methodologies that cater for a variety of negotiating agent architectures. In this paper, we have taken a step towards enabling more effective design of negotiating agents, in which negotiation capabilities are defined in a modular fashion, allowing for reuse, dynamic reconfiguration, and recovery. Furthermore, initial implementations of the construction model have shown that the practical *development process* is also greatly aided by the ability to separate concerns along the different agent views.

This paper sets up the ground for a further investigation of the formal specification and verification of negotiating agent architectures through appropriate conceptual models as well as their empirical analysis through the direct implementation of such models. As such, it brings closer the possibility of negotiating-agent *construction toolkits*, which allow designers to easily configure their negotiators based on application requirements.

References

1. C. Bartolini, C. Preist, and N. R. Jennings. Architecting for reuse: A software framework for automated negotiation. In *Proc. of AOSE-02*, pages 87–98, 2002.
2. M. d’Inverno and M. Luck. Engineering AgentSpeak(L): A Formal Computational Model. *Journal of Logic and Computation*, 8(3):233–260, 1998.
3. M. d’Inverno and M. Luck. *Understanding Agent Systems*. Springer-Verlag, 2001.
4. P. Faratin. *Automated Service Negotiation Between Autonomous Computational Agents*. PhD thesis, UCL, Queen Mary and Westfield, Dept. of Electronic Engineering, 2000.
5. S. Fatima, M. Wooldridge, and N. R. Jennings. Multi-issue negotiation under time constraints. In C. Castelfranchi and L. Johnson, editors, *Proc. of AAMAS-02*, pages 143–150. ACM Press, 2002.
6. N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge. Automated negotiation: prospects, methods and challenges. *Int. Journal of Group Decision and Negotiation*, 10(2):199–215, 2001.
7. N. R. Jennings, S. Parsons, P. Noriega, and C. Sierra. On argumentation-based negotiation. In *Proc. of the Int. Workshop on Multi-Agent Systems*, pages 1–7, Boston, USA, 1998.
8. S. Parsons, C. Sierra, and N. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, 1998.
9. H. Prakken and G. Vreeswijk. Logics for defeasible argumentation. In D. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic*, volume 4, pages 219–318. Kluwer, 2nd edition, 2002.
10. I. Rahwan, L. Sonenberg, and F. Dignum. Towards interest-based negotiation. In *Proc. of AAMAS-03 (to appear)*, Melbourne, Australia, 2003.
11. C. Sierra, N. R. Jennings, P. Noriega, and S. Parsons. A framework for argumentation-based negotiation. In M. Singh, A. Rao, and M. Wooldridge, editors, *Intelligent Agent IV: Proc. of ATAL 1997*, volume 1365 of *LNCS*, pages 177–192. Springer, 1998.
12. TAC. *The Trading Agent Competition*. World Wide Web, <http://www.sics.se/tac/>, 2003.
13. D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Press, Albany, NY, USA, 1995.