Characterizing Multi-Agent Negotiation

H. Van Dyke Parunak
Center for Electronic Commerce, ERIM
PO Box 134001
Ann Arbor, MI 48113-4001 USA
vparunak@erim.org

Abstract

"Negotiation" is a vague term, covering a wide range of interaction mechanisms. This paper resulted from a working group at IWMAS'98.[1] It exhibits a selection of mechanisms that are commonly considered "negotiation," and then develops characterizations for distinguishing different kinds of negotiation, along three dimensions: context management, complexity measurement, and coherence mechanisms.

1 Introduction

In spite of the pervasiveness of negotiation in agent-based systems (or perhaps because it is so pervasive), it is difficult to come up with a formal definition that makes useful distinctions. After several attempts, participants in IWMAS'98 concluded that the benefits of distinguishing negotiation from non-negotiation do not justify the hair-splitting effort involved, and were content to understand the term in a general sense as interaction of any form among agents.

Though it may be academic to rule on whether a particular approach is or is not negotiation, from an implementation perspective it is important to distinguish different approaches from one another, so that we can develop engineering guidelines for matching problem requirements to the appropriate techniques. A working group at IWMAS'98 approached this challenge first by brainstorming a wide variety of negotiation scenarios, and then by reflecting on different dimensions along which they can be compared with one another.

Section 2 presents (without any claim to exhaustiveness) our collection of scenarios. Section 3 outlines three dimensions that we found useful in distinguishing these from one another, again without any claim to exhaustiveness or finality. Section 3.1 discusses the relation between a negotiation and its context, the larger process in which it is embedded. Section 3.2 reviews different mechanisms that agents can use to maintain coherence through the course of a negotiation. Section 3.3 considers different ways in which one can measure the relative complexity of different mechanisms. Each of these sections concludes with implications of the dimension under discussion for implementation. Section 4 offers a summarizing conclusion.

2 Some Examples of Negotiating Systems

Half an hour's brainstorming produced a wide variety of negotiating systems, and was terminated in order to move on to an analysis of characteristics, not for any lack of ideas. This summary offers an *ad hoc* clustering of the systems and applications nominated.

---

[1] The discussion group at IWMAS on this topic included Camille D'Annunzio, Keith Decker, Abhi Deshmukh, Carl Manning, Sandip Sen, Milind Tambe, and Regis Vincent. Some of the group members commented on this paper, but this formulation is not intended to represent a consensus view.

In non-technical contexts, "negotiation" is commonly used in economic and political contexts. Classical examples of economic negotiation include reaching agreement on a contract between management and labor, or on a purchase agreement between buyer and seller. Political negotiations are exemplified by recent episodes concerning the Levant or the Balkans.

Some examples illustrate limiting cases. Perhaps the most complex examples are instances of metanegotiation, in which participants negotiate the conditions under which a subsequent negotiation will take place. These conditions can include the agenda, the participants, or the protocol, and are familiar in the context of middle-eastern peace negotiations. During the main workshop, Jeff Rosenschein offered a useful example from preparations for the 1991 Arab-Israeli talks. The Arabs wanted a Palestinian from Jerusalem at the table, a position at variance with Israel's insistence on the inalienable Jewish status of Jerusalem. The metanegotiators solved the problem by using a boundary definition of Jerusalem accepted by Jordan before the 1967 war brought Jerusalem under Israeli control. This boundary included territory not included within the current Israeli definition of the city limits. A Palestinian representative from this ambiguous territory satisfied Palestinian demands for a Jerusalem representative without compromising the Israeli position.

At the opposite extreme is the pursuit problem, originally introduced to the agent community by Miro Benda at the 1986 DAI Workshop at Twin Lights Manor [2]. Predator agents (e.g., wolves) pursue a prey (e.g., a moose), and succeed only if they can completely surround it. For several years this problem attracted a variety of research approaches drawing from a wide range of sophisticated AI techniques. Korf [4] offers a review of previous work, and reports a simple reactive mechanism that outperforms previous more sophisticated approaches. Korf's predators do not communicate explicitly at all, but avoid one another while moving toward the prey. Perhaps even simpler is the model of a vending machine as a negotiating agent that accepts a piece of currency, determines its genuineness and denomination, and releases a can of beverage if its conditions are satisfied.

Many interesting cases of negotiation are characterized by multiplicities of various sorts. It is common to find *multiple participants*. Air traffic management requires finding an agreement that satisfies the requirements of all aircraft in a region of airspace. The design of a manufactured product requires the agreement of multiple departments, including marketing, service, and manufacturing. Negotiation is often called upon to reconcile *multiple objectives*. Sometimes these are represented by multiple participants in the interchange, as in the case of product design. Service negotiations among telecommunications nodes pose some interesting problems of feature interaction. If a system offers call blocking and also accepts forwarded calls, should it accept or reject a call from a blocked party that has been forwarded through a non-blocked number? If a system permits forwarding, should it automatically forward calls to long-distance numbers? Many cases involve *multiple cycles* to resolve a single issue. Classic iterative auctions (both English and Dutch) fall into this category, in contrast with closed-bid procurement mechanisms that permit only a single cycle.

A specific case of multiple cycles arises in cases that support counterproposals. In assigning weapons to targets, weapons may change the targets they propose to address based on the proposals already on the table from other weapons. In product design, the engineer responsible for one component may propose a design incompatible with the design proposed for a mating part, requiring the designer of the original part to revise its design. Net-based retail sales channels are making it increasingly easy for buyers to engage multiple sales channels in competitive bidding against one another.

It may prove important to distinguish cases with discrete outcomes from those with continuum outcomes. Discrete decisions include cases such as deciding the vendor from which to buy a product, or the weapon to assign to a given target. Continuum decisions are exemplified by family discussions about allocating investments across different classes of securities, or setting the time for supper.

3 Organizing Dimensions

This brief review of cases of negotiation is by no means exhaustive, but does suggest some common dimensions across which different cases may usefully be distinguished. What challenges do agents face in managing their interactions in the context of their environment and the larger processes in which they participate? What mechanisms can they use to maintain coherence in the face of such challenges? How may the complexity of negotiation be measured? These dimensions are not the only way to decompose the space of negotiation, but they do distinguish cases such as those we collected in ways that offer useful guidance to implementers.

3.1 Context Management

Negotiation takes place within a context, defined by the environment in which agents are embedded and the larger processes in which they participate. Negotiation mechanisms differ in how they support a "*going concern*"[2] in which they may be embedded, and in how they respond to threats from the environment.

3.1.1 Relation to a "Going Concern"

Classical AI, and the behavior of many individual agents, is governed by the problem-solving paradigm. A major theme of the workshop was the distinction between this paradigm and the notion of a "going concern." Problem solvers seek achievement goals, and stop when they reach them. A going concerns hopes never to stop, but to maintain itself in perpetual motion in a way that satisfies a set of maintenance goals. A going concern is not a single problem to be solved, but has many facets. (One facet may require problem solving at some time or another.) An important part of the context for a specific negotiation is the going concern in which it is embedded. The working group identified three ways to characterize negotiations in terms of their parent "going concern."

First, negotiation techniques differ in the degree of "common knowledge" about the context that they presume. Cooperative negotiations such as product design are often structured to maximize the amount of common knowledge about the corporate goal being pursued. In competitive negotiations, opponents often seek to conceal their individual objectives from one another. The

[2] The term "going concern is taken from both sociology and the financial industry. In the financial industry, for a business to be a going concern it must be active and operating. The term "going concern" was first used in the context of multi-agent systems by Carl Hewitt of MIT in the introduction of "What is Negotiation" at IWMAS '98. The discussion stemmed from his research on Computer Agencies with Assad Moini and Camille D'Annunzio of The Software Productivity Consortium. The IWMAS'98 video presentations on the web at http://iwmas.wvlink.com/Negotiate/NegotiateHome.nsf?OpenDatabase provide more information

amount of common knowledge must be calibrated in terms of the complexity of the embedding context. The "going concern" involved in middle-east peace talks contains far more information that might be relevant as common knowledge than the "going concern" of a wolf pack seeking to keep itself fed.

Second, negotiation techniques differ in the side effects they impose on the context. One design proposal may be replaced by another without materially changing the course of the negotiation, but a proposal in a peace negotiation or labor discussions can significantly affect the subsequent direction of the talks, whether or not the proposal is accepted.

Third, episodes of negotiation may be linked with one another and with other actions in different ways to support the larger context. An agent may initiate negotiations on a point that is not currently disputed, to preempt anticipated future disagreement or to set the terms of future discussion. Environmental disruptions often require agents to retract commitments from previous negotiations and renegotiate new commitments. In some cases, negotiations yield not only commitments for future action but also sanctions for non-performance, requiring a linkage to mechanisms for enforcement.

3.1.2 Robustness to Threats

Supporting a "going concern" is a positive perspective on the relation of negotiation to the larger context. That context also poses threats that can frustrate negotiation, and negotiation mechanisms differ in the various classes of threats against which they offer robustness. Comparatively simple check-sum and acknowledgement mechanisms suffice to detect and recover from c*ommunication failure*, and in many cases these mechanisms are embedded in the underlying communication system and function invisibly to the application implementer. At the other extreme, it can be difficult to detect *misunderstanding*, when agents successfully exchange messages but honestly interpret them in different ways. Sometimes agents do use the same semantics, but deliberately misrepresent the world. Negotiation mechanisms with penalties for defaulting (e.g., Sandholm's leveled commitment protocol [8]) are intended to cope with the problem of *insincerity* (misrepresentation of commitment), while other mechanisms (e.g., [7]) have been devised that remove the benefits of *duplicity* (misrepresentation of facts). Metalevel negotiation is often invoked to deal with *dynamic change of context*.

3.1.3 Implementation Implications

Negotiation mechanisms are often designed with a focus on the immediate question or task on which agreement between agents is sought. Within such a limited focus, a simple question-answer or contract net protocol may seem sufficient, but is likely to fail in practical application. The discussion in this section shows two further requirements that system designers should address. First, the question or task that is the subject of negotiation is usually embedded in a going concern, and the success of the negotiation will depend on how well it maintains that going concern, not just the resolution of the immediate task at hand. Second, the context can pose a variety of threats to frustrate negotiation, and designers need to review the likelihood of these threats and include appropriate mechanisms to counter them. Awareness of the context within which negotiation takes place invites us to consider negotiation as a situated activity, and to pay special attention to the situating context in designing and implementing a particular protocol.

3.2 Coherence Mechanisms

The dimension of context management classifies instances of negotiation on the basis of the broader contextual objectives of the interaction. This dimension examines specific mechanisms that agents can use to ensure coherence in pursuit of such objectives. Some mechanisms rely only on structures in individual participants, while others depend on structures shared among participants. Mechanisms from these two classes work together to achieve the coherence of an individual negotiation.

3.2.1 Solipsistic Mechanisms

An individual agent can maintain coherence over its end of a negotiation if it has a goal that it is pursuing, a way to measure its current state in comparison with the goal, and a way to plan its next utterance to move closer to the goal. This last requirement inevitably requires some communal mechanism, such as those reviewed in the next section.

Negotiations may be distinguished by the classes of goals in pursuit of which they are useful. An extensive typology is offered in [11]. The working group enumerated the following distinctions.

- Relatively simple and bounded negotiations are often sufficient to support *static goals*. *Changing goals* introduce the need for renegotiation or meta-level negotiation.
- Mechanisms guarding against malicious threats (such as duplicity or insincerity) are more necessary when an agent's goals are *Contentious* than when its intents are *Cooperative*.
- An agent will manage its end of a negotiation differently depending on whether the domain over which its utility function is defined is *Discrete* or *Continuous*. In particular, different classes of optimization algorithms may be necessary to evaluate the current state of the negotiation and plan the next step.
- Most real-world problems must satisfy *multiple objectives* concurrently, a requirement that is considerably more difficult than *unidimensional* goals. One of the values of market mechanisms is the potential for reducing multiple objectives to a common currency, thus turning multidimensional problems into unidimensional ones.
- An agent negotiating over *future action* (commitments) must address issues of time management that do not arise in negotiations over *present action*.

3.2.2 Communal Mechanisms

In addition to the individual mechanisms considered in the previous section, coherence of a negotiation requires that the participants have certain things in common. This "common knowledge" most commonly takes the form of a common language, but may consist only of the laws of physics in a shared environment.

A *shared language* may exist at a number of levels of sophistication.

- Perhaps the simplest is a common *ontology*, an agreement on how to divide up the world, supported by a shared *vocabulary* and (in a market system) a *currency* with a generally agreed-upon valuation.

- In addition to the vocabulary to use within utterances, agreement is usually needed on the *metalevel classification* of the kinds of utterances that will be supported, such as rudimentary types of speech acts [1, 9] or structured templates for various classes of discourse [3, 10].
- A *static protocol* (such as the contract net) outlines a fixed order of utterances of various categories.
- More sophisticated negotiation is possible when agents can reason about the stream of speech acts they have received, and decide dynamically what sorts of utterances are appropriate.

Even in the absence of a common artificial language, agents can interact by changing a *shared physical environment* and sensing those changes. Wolves who surround a moose by avoiding each other while moving toward the moose [4] exemplify this sort of minimal reactive coordination. Though insect pheromones are not a naturally-occurring part of the environment, the response of insects to these fields is most likely hardwired rather than reasoned, and thus may be usefully be considered another instance of environmental reactivity. This latter case exemplifies the potential of flow fields (such as a pheromone gradient) to organize and coordinate the movements of multiple agents [5, 6].

3.2.3 Implementation Implications

There is no single best mechanism for maintaining coherence through a single negotiation or across a sequence of negotiations in support of a going concern. A classification of coherence mechanisms such as this is one foundation for constructing a body of engineering knowledge matching different mechanisms with the classes of problems for which they are most appropriate. Another is outlined in the next section.

3.3 Complexity Measurement

Perhaps the most straightforward way to compare instances of negotiation is on the basis of their complexity. Yet in some ways this is the least important dimension for comparison. Negotiation exists only because it needs to achieve an objective. Simplicity is no virtue if the task at hand is not accomplished, nor is complexity if it unnecessarily increases the cost and bulk of a solution. Complexity measures should be applied both to negotiation mechanisms and to the problems to which they are applied, with the objective of developing engineering guidelines for selecting appropriate mechanisms for appropriate problems.

3.3.1 Complexity of Mechanisms

The complexity of a negotiation mechanism may be assessed at the level of the community, the protocol, or the overall process.

One can assess the complexity of the agent *community* in which negotiation takes place. If one views the community as a graph in which nodes are agents and edges represent the possibility of a message between two agents in the course of the negotiation, one can invoke a variety of graph-theoretic measures, such as the total number of agents, the number of components in the graph, the population and diameter of each component, and distributions over the degrees of the various nodes. It is also useful to consider the complexity of individual agents (for instance, reactive vs. BDI), or the degree of knowledge that must be shared across participants to make the

Table 1: A Cline of Protocols

Class of Protocol	Description	Minimal Language
Reactive	Sense, then Act	Environment
Command	Master agent sends unilateral instructions to servants	Symbolic
Voting	One-shot quantitative statement of interest	Currency
Fixed Protocol	Back-and-forth; symbolic & quantitative	Messages & Protocol (Contract Net)
Conversation	Arbitrary interchange	Speech Acts (statement, request, commitment; command; e.g., KQML)

negotiation succeed (for example, possession of a common ontology vs. knowledge of one another's utility functions).

The *protocol* that governs interactions among participants is a specialization of the notion of "degree of shared knowledge" that deserves discussion as a separate category. Table 1 exemplifies several classes of protocol arranged in a cline from less complex to more complex.

The classification of protocols in the previous paragraph is static, but what matters for practical application is the dynamic complexity of the resulting *process*. If there is a well-defined goal for the negotiation (as in a problem-solving community), one may be able to characterize the temporal complexity of the negotiation process and compare it with a serial process for attaining the same goal in a single-threaded reasoner. When a formal characterization of the process is not available, an external observer may assess the complexity of the process in less formal terms, but such an assessment will be a function of the observer's complexity as well as the intrinsic complexity of the process. However assessed, process complexity will depend not only on the intrinsic negotiation mechanisms being used, but also on the complexity of the problem being addressed.

3.3.2 Complexity of the Problem

Assessing the complexity of the problem being addressed is important for two reasons. First, it helps determine the degree to which observed process complexity is a function of the intrinsic mechanisms being used. Some mechanisms may yield simple processes on some problems, but complex ones on others. Second, measures of the complexity of a specific problem can be a useful engineering guide to the appropriate negotiation mechanisms to apply to that problem. Useful problem measures include the size and topology of the problem space, and the size and topology of the subspace that constitutes the desired solution.

3.3.3 Implementation Implications

All things being equal, good engineering strives for simplicity, and agent implementers should prefer the simplest systems that will do the job. Simplicity has many benefits, including shorter development time, systems that are easier for users to understand and maintain, and reduced platform costs. However, it can be difficult to tell when all things are equal. Complexity measures interact with one another in ways that directly impact implementation decisions. For example, there is often a trade-off between process complexity and protocol complexity, such that more flexible protocols may lead to faster convergence than simpler ones (or even achieve convergence in cases that are intractable for simple protocols). The implementer must assess the degree of this interaction in deciding the impact of agent simplicity and speed of convergence in a given application. Similarly, if problem complexity can be bounded, a simpler protocol may achieve satisfactory results at lower cost than if one attempts to construct a general-purpose agent that is over-engineered for the problem at hand.

4 Conclusion

Discussions at IWMAS'98 revealed that "negotiation" is a broad and amorphous topic, including almost every form of interaction among computational processes that can be imagined. Scientific progress in *studying* negotiation requires a way for researchers to characterize the object of their study, so that conclusions from different efforts can be compared with one another. Engineering progress in *using* negotiation requires a disciplined way to characterize different mechanisms as a first step to defining the conditions under which each is an appropriate implementation choice. The working group discussion captured in this document does not yield the formal ontology one might desire, but does suggest some preliminary categories that may be useful to both scientists and engineers in explaining their work to one another and in identifying under-explored topics for further investigation.

References

[1] J. L. Austin. *How to Do Things with Words*. Oxford University Press, 1962.

[2] M. Benda, V. Jagannathan, and R. Dodhiawalla. On optimal cooperation of knowledge sources. Boeing Computer Services, Bellevue, WA, 1985.

[3] J. Conklin and M. L. Begeman. gIBIS: A Hypertext Tool for Exploratory Policy Discussion. MCC, Austin, TX, 1988.

[4] R. E. Korf. A Simple Solution to Pursuit Games. In *Proceedings of Eleventh International Workshop on Distributed Artificial Intelligence*, 1992.

[5] P. N. Kugler and M. T. Turvey. *Information, Natural Law, and the Self-Assembly of Rhythmic Movement*. Lawrence Erlbaum, 1987.

[6] H. V. D. Parunak. 'Go to the Ant': Engineering Principles from Natural Agent Systems. *Annals of Operations Research*, 75:69-101, 1997. Available at http://www.erim.org/~van/gotoant.pdf.

[7] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. Cambridge, MA, MIT Press, 1994.

[8] T. W. Sandholm. *Negotiation among Self-Interested Computationally Limited Agents*. University of Massachusetts, Computer Science, 1996.

[9] I. A. Smith and P. R. Cohen. Toward a Semantics for a Speech Act Based Agent Communications Language. In *Proceedings of CIKM '95 Workshop on Intelligent Information Agents*, 1995.

[10] S. Toulmin. *The Uses of Argument*. Cambridge, UK, Cambridge University Press, 1958.

[11] D. N. Walton. *Informal logic : a handbook for critical argumentation*. New York, NY, Cambridge University Press, 1989.