

NEGOTIATED DECOMMITMENT IN A COLLABORATIVE AGENT
ENVIRONMENT

by

Michele Miley Van Dyne
B.A., University of Montana, 1981
M.S. University of Montana, 1985

Submitted to the Department of
Electrical Engineering and Computer Science
and the Faculty of
the Graduate School of the University
of Kansas in partial fulfillment of
the requirements for the degree of
Doctor of Philosophy

Dr. Costas Tsatsoulis, Committee Chair

Dr. W. Perry Alexander, Committee Member

Dr. John Gauch, Committee Member

Dr. Susan Gauch, Committee Member

Dr. Prakash Shenoy, Committee Member

Date Submitted: _____

ABSTRACT

Michele Miley Van Dyne, Ph.D.

Department of Electrical Engineering and Computer Science, May 2003

University of Kansas

Decommitment of agents has been studied in the context of self-interested agent societies, primarily using a leveled commitment protocol in which agents committing a breach of contract are penalized in various ways. Limited research has investigated decommitment in collaborative agent populations. This research investigates decommitment in a collaborative and benevolent agent society, in which agents make decisions and act in a manner which maximizes the global benefit. In this environment, decommitment takes on a different nuance than much of what has been presented in the literature, in that collaborative agents can be expected to behave so as not to cause harm to the agent society as a whole, even in decisions not to honor commitments.

When an agent perceives that decommitting to a previous commitment has the potential for increasing its local utility, the agent first attempts to negotiate the terms of the decommitment with those affected. In the context of collaborative agents, it was hypothesized that allowing agents to decommit would have an overall beneficial effect on attaining system goals, however, decommitment would require additional overhead as far as information required and communication resources. Thus, there would be a trade-off between the resource constraints placed on the system versus the beneficial effect of decommitment. Negotiated decommitment was expected to provide more of a beneficial effect than unilateral decommitment.

Testing of the theory and experimentation was performed in the context of the Autonomous Negotiating Teams (ANTS) system. Agents are based on the BDI agent architecture model and have the characteristics that they are rational, collaborative, benevolent, homogeneous and multitasking, and they are capable of negotiation and planning time-dependent activities. The operating environment is both time-bounded and resource constrained.

Results indicated that negotiated decommitment performed as hypothesized, that is, overall performance improved. Graceful degradation of performance was demonstrated in all conditions as the constraints on the system increased. Under all conditions, agents demonstrated rationality in decision making and both negotiated decommitment and unilateral decommitment provided a higher locally assessed utility than did the baseline condition.

ACKNOWLEDGEMENTS

I would like to express my sincere appreciation to Dr. Costas Tsatsoulis, my advisor during most of my time at the University of Kansas. I greatly appreciate both the encouragement and correction he provided, and his patience. Without his guidance, support and encouragement, this work would not have been even remotely possible.

I would also like to thank the rest of my committee members: Dr. W. Perry Alexander, Dr. John Gauch, Dr. Susan Gauch, and Dr. Prakash Shenoy. Their advice and guidance on my proposed research was very valuable in defining the problem and in carrying out different aspects of the work.

Finally, I would like to thank Dr. Linda Goodwin of Duke University, who provided much encouragement during the completion of this thesis, but most of all, I appreciate her patience and understanding of the priority of completing this work, at the expense of the timeliness of other obligations.

The work described in this dissertation was sponsored in part by the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory, Air Force Material Command, USAF, under agreement number F30602-99-2-0502. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency (DARPA), the Air Force Research Laboratory, or the U.S. Government.

Table of Contents

Chapter 1 – Introduction	1
1 Research Problem.....	5
1.1 Research Issues.....	5
1.1.1 Intuitive Definition of Commitment and Decommittment.....	6
1.1.2 Why Decommit?	7
1.1.3 Decommittment Repercussions.....	8
1.2 Hypotheses.....	9
Chapter 2 – Background and Related Research.....	10
2 General Background on Agents	10
2.1 Agents.....	10
2.2 Agent Negotiation.....	20
2.3 Coalition Formation.....	29
2.4 Commitment and Decommittment	34
Chapter 3 – Formalization of the Problem.....	39
3 Approach	39
3.1 Distributed Task Scheduling	39
3.2 Individual Agents and Agent Society.....	41
3.2.1 Agent Society and Interaction.....	43
3.3 Negotiation	44
3.3.1 Agent Commitment and Decommittment	46
Chapter 4 – Implementation.....	49
4.1 Problem Domain.....	49

4.2 Agent Software	54
4.2.1 Agent Characteristics	54
4.2.2 Agent Architecture	55
4.2.3 Agent Interaction.....	60
4.2.4 Local Estimate of Global Utility	64
4.2.5 Decision Criteria	68
Chapter 5 – Experimental Design	76
5.1 Performance Evaluation Criteria.....	77
5.2 Experimental Conditions	79
Chapter 6 – Results and Analysis.....	84
5.1 Overall Goal Achievement	84
5.2 Graceful Degradation of Performance.....	93
Chapter 7 – Conclusions and Future Directions.....	102
References	106

List of Figures and Tables

Figure 4.1: Radar Simulator Screen; 6 Sensors, 2 Targets.....	52
Figure 4.2: Agent Software and External Programs	53
Figure 4.3: Agent Multithreading.....	56
Figure 4.4: Communicator Operation	57
Figure 4.5: Scheduler Operation	57
Figure 4.6: Negotiator Operation	59
Figure 4.7: Decision Maker (Agent) Operation	60
Figure 4.8: Request Negotiation Protocol.....	62
Figure 4.9: Negotiated Decommittment.....	64
Table 5.1: Experimental Variables and Level of Constraint.....	80
Figure 5.1: Sensor Position and Orientation with Four Agents	81
Figure 5.2: Sensor Position and Orientation with Six Agents.....	81
Figure 5.3: Sensor Position and Orientation with Eight Agents	81
Figure 5.4: Sensor Position and Orientation with Twelve Agents.....	81
Figure 5.5: Target Position.....	82
Figure 5.6: Curved Target Path: Oval	83
Figure 6.1: Average Number of Planned Measurements Across all Experiments within an Experimental Condition	85
Figure 6.2: Average Number of Times 3 or More Measurements were Taken in a Two Second Window Across all Experiments within an Experimental Condition	86
Figure 6.3: Balance of Measurements across Multiple Targets	87

Figure 6.4: Average Total Number of Measurements Across all Experiments within an Experimental Condition	88
Figure 6.5: Average Tracking Error Across all Experiments within an Experimental Condition.....	89
Table 6.1: Average Utility per Millisecond	90
Figure 6.6: Average Tracking Error Across Levels of Constraint	94
Figure 6.7: Average Tracking Error Across All Experimental Conditions by Target Speed	95
Figure 6.8: Average Tracking Error Across all Experiments by Target to Sensor Ratio	96
Figure 6.9: Average Planned Measurements by Experimental Condition.....	97
Figure 6.10: Average Times Three or More Measurements were Taken in a Two Second Window	98
Figure 6.11: Averaged Standard Deviation of Measures Across Multiple Targets	99
Figure 6.12: Average Total Number of Measurements per Target	100
Figure 6.13: Average Tracking Error Across Increasing Constraints.....	101

Chapter 1 – Introduction

In a society of collaborative and benevolent agents, it seems reasonable to expect that if an agent makes a commitment, then that commitment will be honored. A more realistic view, however, is that there may be conditions under which that commitment will not be honored, and in some cases, it is preferable not to honor a commitment when circumstances change.

On a basic level, the issue of agent decommitment has been addressed in the Belief-Desire-Intention (BDI) agent architecture. Rao and Georgeff (1991, 1995) describe three types of agent behavior dependent on what types of commitment termination conditions are allowed by a BDI agent. “Blindly-committed” agents will deny any changes to beliefs or desires that conflict with their commitments. A “single-minded” agent will allow changes to its beliefs and may drop commitments accordingly. Last, an “open-minded” agent will allow changes in both beliefs and desires that may force commitments to be dropped. This addresses characteristics an individual agent may possess that may lead to decommitting to a previous commitment, but it disregards decommitment effects on other aspects of the system.

Schut and Woolridge (2001) extend the BDI model to include a decision theoretic approach to what they call “intention reconsideration.” Their approach allows an

agent to decide for itself (as opposed to having an a priori commitment level) whether to perform an action or to deliberate over its intentions. Two conditions are defined under which reconsideration may be appropriate: 1) when conditions have changed such that an intention can no longer be achieved, it may be good to drop that intention, and 2) when opportunities arise it may be desirable to adopt additional intentions. While this approach moves the decommitment decisions from design-time to runtime, it only allows an agent to consider what action should be done “next”, a default action or a reconsideration. The approach also studied only the effects on a single-agent system, and thus still disregards effects on the system as a whole in a multi-agent system.

Norman, Sierra and Jennings (1998) have defined a formalization of agent commitment. In their view, an agent must be capable of performing an action and have the right to perform that action, before it commits itself to the task. But, once it agrees to perform the act, it is “bound” to do it, or committed. They discuss the fact that their approach does not account for decommitment, or “loss of commitment,” which they say may happen because the action is no longer possible, the agents have agreed to be no longer bound to that agreement, or because the agreement has expired. Andersson and Sandholm (1998) mention conditional commitment as one way to address some decommitment actions, as it expresses a less firm notion of commitment. Conditional commitment has been expressed in “contingency contracts” (Raiffa, 1982) with other agents, in which an agent may commit to a request on the

condition that some future event occur, however, Andersson and Sandholm (1998) state that it may be impossible to anticipate all future events, and thus conditional commitments can only cover a portion of all commitments that may be broken.

Excelente-Toledo, Bourne and Jennings (2001) introduced the concept of differing degrees of commitment, along with differing types of sanctions for decommitment, in a population of rational and self-interested agents. Experimentally, they showed that an intermediate degree of commitment, which implies a certain degree of loyalty, leads to better performance than either total commitment or loose commitment.

Setting imposed penalties based on the prevailing context at the time of decommitment provided better performance than either fixed or partially sanctioned penalties. Because the agent population was self-interested, the goal of each agent was to maximize its own gain, however, which implies the concept does not necessarily directly extend to a fully cooperative agent population, in which the goal of each individual agent is to further the global goal of the entire system.

The contract net protocol (Smith, 1980) is a well known approach to agent negotiation, in which a requesting agent acts as a manager and requests bids from other agents (contractors) to perform a task. Agents respond to this request by submitting bids and the requestor chooses from among those bidders. Xing and Singh (2001) look at the role of commitment under the contract net protocol type of agent interaction and formalize their approach to both commitments and decommitment. In

this research, agents are self-interested, and have certain conditions under which they may decommit. No penalties were discussed. Sandholm and Lesser (1996) discuss the use of a “leveled commitment contracting protocol” in conjunction with the contract net protocol. Again, agents are self-interested, and any agent may decommit unilaterally, based on its own local reasoning. A penalty must be paid for decommitment. When overall system performance is measured, it shows improvement, but agent decommitment decisions are made based on analysis of their own individual gain.

Sen and Durfee (1996) discuss the use of decommitment, or cancellation as it is termed in their work, to resolve conflicts in a meeting scheduling domain. In their discussion, cancellation occurs when another meeting of higher priority can only be scheduled at a time when a lower priority meeting has already been scheduled. In this approach, the cancellation of a single member of the meeting would cause the cancellation and subsequent rescheduling of that meeting. Resolving conflicts through cancellation and rescheduling was left as an initial investigation in which they discuss some of the issues surrounding what they term a “significant problem”, and was not experimentally evaluated or validated in their work.

The domain area used for testing in this research is that of multi-sensor target tracking in which each agent has one sensor (radar) with three sectors that allow it a 360° view of a distance-limited range. Only one sector may be active at any given time. As a

target moves through the area, agents controlling sensors in different locations must cooperate to track the target as it moves in and out of their sensing ranges. While the global problem is one of tracking a target through space, each agent has only a local view of the situation and therefore must negotiate with others to perform tasks it cannot perform alone. Furthermore, a moving target implies that the problem is time-dependent, so agents must negotiate and act within a reasonable amount of time, relative to the target speed.

1 Research Problem

1.1 Research Issues

This research investigates the decommitment of agents in a collaborative and benevolent society of agents. Andersson and Sandholm (1998) call these “social welfare” or SW agents, in which each agent attempts to maximize the profit of the whole society, as opposed to self-interested agents that attempt to maximize their own individual profit. In the society of interest to this research, agents are multitasking, such that they can perform certain actions in parallel, and they have an agenda, or schedule, of what tasks they intend on performing at certain time-stamped points in the future, such that they are not constrained to only contemplate what action to perform “next”. Agents operate in a resource-bounded environment, and actions are constrained by the time at which they must be performed. In order to request that another agent perform a task, agents are capable of negotiation.

The entities affected by a decommitment may vary. Decommitment can occur under differing conditions, for different reasons, both intentional and accidental.

Furthermore, different circumstances of decommitment have different effects on the system as a whole. These issues are discussed in the following sections.

1.1.1 Intuitive Definition of Commitment and Decommitment

When an agent decides to perform a task, we can view that as a commitment. In the BDI model of agent architecture (Rao and Georgeff, 1991, 1995), this can roughly be viewed as an “intention.” An agent may intend to perform an action because it has its own goals to attain (“desires” in the BDI framework) and the information available to it (“beliefs”) indicates that it is a feasible thing to do, or because another agent has requested that it perform that task. Thus, we have task commitments to either oneself or to another agent. In the terminology of the BDI model, task commitments are “intentions to,” and indicate the commitment to perform a particular action.

Decommitment may be thought of as renegeing on one of these commitments.

Decommitment can be either intentional (rational) or accidental. An agent may make a conscious decision not to live up to a commitment, or it can attempt to meet a commitment and fail. Note that, in this research, decommitment is considered to be a rational act performed by the committed agent. Should an agent find that it is not

physically possible to perform a task it had committed to, this is not considered a decommitment for the purposes of this investigation.

1.1.2 Why Decommit?

There are several circumstances under which an agent may decommit. An agent may find that it is physically impossible to perform the task, and thus it may be forced to decommit. The agent may also be presented with an alternate task that has a higher degree of immediacy or importance and it may choose to perform that action if it believes it will benefit the society more than its original committed task would have. Finally, if conditions indicate that the performance of a task is no longer a productive activity, the agent may decide not to honor its commitment. The last two conditions indicate rational decisions on the agent's part to decommit, while the first condition is one in which decommitment is unavoidable. For the purposes of this research, only intentional decommitments are investigated, since unavoidable decommitments are non-negotiable.

More specifically, then, an agent will consider decommitment when it attempts to schedule a task and finds that another task of lesser importance conflicts with the new potential commitment.

1.1.3 Decommitment Repercussions

Just as commitments in an agent society have an effect on the overall performance of the society, presumably beneficial or they would not be undertaken, decommitments also have an effect. In a collaborative society, where actions of the individual agents are chosen only by their positive effect on the society as a whole, however, we can assume that voluntary decommitments will only be performed when an agent perceives that it is in the best interest of the group, or that it is rational, to do so. Even voluntary decommitments which increase overall social welfare, may be detrimental to some subset of the agent system, though. Thus, we need to consider the impact of decommitment.

The *value* of a commitment is simply an agent's local estimate of the global utility of performing that action. When a commitment is made, the value is either determined through parameters passed via negotiation with another agent, and by information the agent itself has sensed about the environment. Since the environment is dynamic, the value of a commitment may change over time. The value will be recalculated if another task of seemingly higher value presents itself.

The second aspect of a commitment is its *strength*. The strength of a commitment is determined by the agent's perception of the impact it would have if broken. A commitment is stronger if the time to perform the committed action is near or if the

commitment was made by request of another agent. In other words, strength is a measure of the impact of dropping a commitment.

Both value and strength must be taken into consideration before an agent decides to drop a commitment, and together they form an agent's local assessment of estimated global utility.

1.2 Hypotheses

Given a society of agents that are characterized as autonomous, rational, collaborative, benevolent, and that are capable of negotiation and planning time-dependent activities:

1. Allowing agent decommitment will improve the overall goal achievement of the system.
2. When an agent's decommitment to a task or goal affects other agents, negotiating the terms of the decommitment with those other agents will be more beneficial than unilateral decommitment, despite the associated increase in overhead. The improvement is expected as a result of a more refined estimation of global utility when more than one agent contributes to that estimate.
3. As the number of constraints on the system increase, the overall goal achievement of the system will degrade gracefully.

Chapter 2 – Background and Related Research

2 General Background on Agents

2.1 Agents

Rao and Georgeff (1991, 1995) describe individual agents from both a theoretical and practical standpoint by defining them in terms of Belief-Desire-Intention (BDI) logics. BDI agents have the mental attitudes of belief (information), desire (motivation), and intention (deliberateness). Rao and Georgeff argue that real-time application domains have characteristics that require the use of BDI since: 1) the environment is non-deterministic, 2) the system itself is non-deterministic, 3) at any point in time, there may be many objectives for the system, 4) the best action to achieve an objective is dependent on the environment and not the state of the system, 5) the environment can only be sensed locally, and 6) actions can be carried out within reasonable bounds to the rate at which the environment changes. Given these characteristics, the authors claim a decision-theoretic approach cannot be used, particularly because of the rate at which the environment may change. Rao and Georgeff describe the process of moving from a decision-theoretic model to the BDI model, and then provide an architecture for implementing an agent systems using BDI. Rao (1995) describes an extension to the BDI model for recognizing the mental

states of other agents, which is particularly useful in adversarial agent systems, and potentially helpful in cooperative yet heterogeneous systems.

Woolridge and Jennings (1995) describe a “weak notion of agency” as having the following four characteristics: autonomy, social ability, reactivity, and pro-activeness. Autonomous agents have some degree of control over their actions and internal state without the direct intervention of others. Agents possess a means of communication between other agents and humans, giving them the characteristic of social ability. Agents perceive and respond to their environment and changes in that environment in some manner, meaning they are reactive. Finally, they are pro-active in the sense that they exhibit goal directed behavior and thus do not only respond to changes in their environment.

In earlier work, Decker (1987) took a different approach in examining agent system properties, rather than individual agents, by looking at the dimensions of control and communications as the primary characteristics. Control is comprised of the attributes of the degree of cooperation between agents, the organizational structure of the agents, and how cooperation is achieved (which he terms “dynamics”). The cooperation dimension may run the gamut from fully cooperative agents to antagonistic agents. Agent organization may be anywhere from totally structured or hierarchical in some form, to a completely peer to peer group of agents with no pre-ordained structure. The dimension of dynamics, or how cooperation is achieved, can

vary from statically determined prior to run time or dynamically determined dependent on the conditions encountered at run time. Decker divides the communication dimension into the sub-dimensions of protocol, content, and paradigm. Paradigm refers to the medium through which agents communicate, whether via a global memory or through message passing, or some combination of the two. The content of a communication may be completely accurate or it may contain uncertain or incomplete information. Finally, protocol refers to the high-level (or application layer) means of communication as opposed to the low-level mechanics. Messages may be sent as selective or broadcast, a characteristic that describes who receives a message. They may be either on-demand (an agent requests that the message be sent), or they may be unsolicited. Acknowledgement by the receiving agent may be part of the protocol or not, which is one method of handling errors. Decker states that the higher the degree of cooperation between agents, the more potential for communications. However, it is desirable to minimize the amount of communication simply because communication bandwidth is most often a limited resource.

Jennings (2000) emphasizes the social nature of agent systems. He addresses agent-based software engineering and includes these interaction and organizational aspects in his discussion. First of all, he discusses the use of methods designed to decrease software complexity by Booch (1994) to aid in software engineering in agent-based systems. These methods are decomposition, abstraction and organization, all three of

with apply directly to the simplification of describing agent architecture. But perhaps most importantly, Jennings proposes the notion of a “social level” to study an agent-based system’s behavior and structure without delving into implementation or interaction protocol details. This social level view is discussed as a next level above Newell’s (1982) knowledge level view of individual asocial intelligent agents. In this social view, instead of studying an asocial agent, the agent organization is the system of interest. The agents, their interaction channels, dependencies and organizational relationships are the components of interest, as opposed to the goal and action components of the knowledge level. Compositional laws include agent roles and organization rules, and the behavior law is a principle of organizational rationality, in contrast to the principle of rationality in the knowledge level view. Finally, knowledge is the medium in the knowledge level view but organization and social obligations, means of influencing others, and means of changing organizational structures are the media at the social level. Jennings suggests that this view may provide a basis for developing tools that will aid in developing agent-based systems.

Castelfranchi (1998) also stressed the importance of social action in agent-based systems. In his view, “AI is the discipline aimed at understanding intelligent beings by constructing intelligent systems”, and intelligence is a social phenomenon, therefore we must build social systems in order to have intelligent systems.

Castelfranchi discusses several concepts that distinguish the view of the agent as an individual versus the social nature of agent systems. First, he discriminates between

social action versus collective action and concludes that we cannot understand social action by merely looking at individual actions of agents unless we also look at the social character of the individual action. Second, agents not only need mental states, such as in the BDI model, but they should be able to represent the mental states of other agents. He terms this criterion “mind reading”. Third, he distinguishes between social action and communication, stating that communication does not mean agents are social, but that agents communicate because they are social. Fourth, social action is not necessarily cooperation between agents. Finally, he discusses reconciling emergent and cognitive behaviors. In his view, cognitive behaviors may be emergent.

In examining social behaviors from the bottom up, Castelfranchi (1998) examines different levels of action. At the base of the action list is the goal-oriented nature of agents and their actions. He specifies that agents must be cognitive in order to have goals, and that their actions are then goal-directed if their actions are directed by goals and if their goals, decisions, and plans are based on their beliefs (in the BDI model, the information agents possess about the environment). At this level, there is not necessarily any interaction between agents. The next level, which Castelfranchi sees as the first step toward social action are the concepts of interference and dependence. Interference does not necessarily imply a negative relationship, just that the effects of an action of one agent may be relevant to the goals of another. With dependence, an agent cannot achieve one of its goals without a resource or action supplied by another agent. Two forms of action may occur at this level. An agent may

adapt its behavior to the actions of another in order to use the other's action to its own benefit or to avoid some sort of negative interference. The other action is to attempt to change the other agent's behavior to either do something helpful or to not do something harmful. The third level of social action (and the second step in social action) is for an agent to have beliefs about another agent's mind. Goals are what distinguish actions from social actions since the exact same behavior may be exhibited with or without considering another agent. At this level, Castelfranchi defines "social beliefs" as beliefs about another agent's mind or actions, and defines this a weak form of social action. These social beliefs relate back to the notion of mind reading mentioned earlier, where an agent has a representation of the goals and beliefs of another agent.

At the next level, though not the next step in the hierarchy of social action, Castelfranchi (1998) discusses coordination of behavior among agents, and distinguishes two types of coordination. The first is that of reactive coordination, which is, as it implies, based on the reaction to a perceived threat or opportunity. The second type of coordination is anticipatory coordination, based on the anticipation, not the perception, of a threat or opportunity. He states that no agent can really plan its behavior without some level of anticipatory coordination. To further discriminate about coordination, Castelfranchi talks about three levels: unilateral, bilateral and mutual coordination. If an agent is only coordinating its own actions with respect to another, that is unilateral. If the other agent happens to be doing the same with respect

to the first agent, then it is bilateral. However, the coordination is mutual if both agents are aware of their coordination activities and reach some sort of agreement concerning them. (In earlier work, a similar taxonomy was applied to what was termed “cooperation” in Conte, Miceli and Castelfranchi (1991), where “accidental cooperation” is analogous to bilateral coordination, “unilaterally intended cooperation” is similar to unilateral coordination, and “mutual cooperation” is analogous to mutual coordination.) To further the concept of coordination, in Castelfranchi (1998), there is also the dimension of selfish versus collaborative coordination. In selfish coordination, an agent is only coordinating with another to further its own goals, but in collaborative coordination, the agent may be coordinating in order to further the goals of the other agent and not its own. The fifth level, and third step of social action, is defined as relying on another agent, or delegation. Delegation is based on dependence on another agent. If the first agent cannot reach its goals without the assistance of another, that is strong dependence, however, if the first agent simply prefers to free its own resources by having another agent perform a task, then it is considered weak dependence. Delegation may be weak, where the first agent simply takes advantage of another agent performing a needed action, it may be somewhat stronger if the first agent induces the second in some way to perform an action, or it may be strong if the second agent is aware of the first agent’s need for the action and has agreed to perform the task. This last form is termed “Social Goal-Adoption”.

The sixth level, and fourth step in social action, described by Castelfranchi is for an agent to have goals about another's actions and goals. That is, the first agent not only has the goal that the second agent perform an action, but it also has the goal that the second agent have a goal to do that action. He terms this "cognitive delegation". In order for an agent to adopt a new goal, that agent must have some change in beliefs that makes the goal appropriate. Therefore, the first agent must have some way of influencing the beliefs of another agent. Furthermore, the first agent must have some sort of power over the second in order to cause this change of beliefs and goals, only one form of which is benevolence of the second agent. The seventh level in Castelfranchi's hierarchy, and fifth step in social action, is that of social goal adoption, defined previously under the topic of delegation. If an agent adopts another agent's goal, it comes to either have a new goal, or to have additional reasons for pursuing a goal it already had. To carry this one step further, if the second agent adopts a goal of the first agent, and also adopts the first agent's goal that it adopts the first agent's goal, it is called "goal adhesion" or "compliance". This is considered the strongest form of goal adoption. Of course, if the second agent is truly autonomous, it may decide on its own whether or not to adopt the first agent's goals. Goal adoption need not be based on benevolence, but may also occur if the second agent perceives it as being instrumental to achieving its own goals, that is, there are selfish motives for adopting another's goals. Another option is that both agents share goals and both have an interest in achieving them, which makes goal adoption a viable option. There are different levels of collaboration once an agent has agreed to adopt a goal. There is

literal help, in which an agent adopts exactly what was asked of it. There is overhelp, in which an agent goes beyond what was asked, without changing the original goal. There is critical help in which an agent satisfies the original request but in some manner changes the plan or action along the way. There is overcritical help where an agent provides critical help as before, but also goes beyond that help. Finally, there is hyper-critical help in which an agent adopts additional goals that the requesting agent didn't originally consider, and in so doing may not fulfill the request as originally stated. Obviously, only cognitive agents can do more than fulfill what was originally asked, in that they need to infer what may truly be needed by the requesting agent.

The final notions of social action proposed by Castelfranchi are that social goals are the "glue of joint action", and that social structures and organizations are very important in social action. In discussing social goals, he defines "social commitment" as the result of strong delegation on the part of the first agent, and strong adoption on the part of the second agent. This goes beyond individual commitment to a goal by an agent, where an agent has a commitment to a goal, which may be abandoned should it decide that the goal has been reached, that it is impossible to reach, or that the agent is no longer motivated to reach that goal. Social commitment is a relational concept in which an agent is committed to another agent, and not just to a goal. Social commitment creates rights and duties between two agents. Social commitment is just one aspect of social organization or structure. Other aspects of social structure are the

interdependence and power structure of the agents, the personal acquaintance structure between agents, the communication structure, the commitment structure, and any structure determined by existing rules and norms about action and interaction. Because of organizational structure, individual agents are not as free to commit themselves as they might like. In addition to the above aspects of structure, a dependence structure may also emerge as agents interact. Mutual dependence may occur when two agents depend on each other in order to realize a common goal, and reciprocal dependence may occur when two agents depend on each other to achieve different goals. According to Castelfranchi, the resulting dependence network will determine and predict partnerships, coalition formation, competition, cooperation, exchange, functional structure in the organization, rational and effective communication, and negotiation power.

Viewing an agent-based system as a social system gives rise to the potential need for social laws. Shoham and Tennenholtz (1992) describe a social law as one that imposes constraints upon the actions that may be programmed into agents. They describe a “useful” social law as one that guarantees an agent still has a legal plan to move between focal states, a subset of the full state space. Finding a useful social law, in the general case was shown to be NP-complete, however under certain restrictions to the system, it can be shown to be polynomial. Fitoussi and Tennenholtz (2000) informally describe a social law as one that restricts the legal actions available to an agent such that conflicts between agents are minimized. They focus on two

types of social laws that are preferable to arbitrarily determined laws: minimal laws and simple laws. A minimal social law is one that minimizes the constraints upon the agents such that individual flexibility is maximized. In their view, such a law is only “useful” if it still allows agents to meet their goals, both for safety to the agent and for whatever purpose the agent was designed (these are called “liveness goals”). A simple social law is one which requires as little of the agent as possible in order to obey it. The necessity of such a criteria is more obvious in domains where agents have limited sensory capability, or under conditions where sensors may fail and the agent then will be unable to follow a more complex law. The same criteria of usefulness applies to simple laws, just as it applies to minimal laws. Minimal laws and simple laws are not necessarily the same thing, dependent upon the domain problem, and deciding whether either a minimal law or a simple law exist for a given system is NP-hard. However, starting with a law, Fitoussi and Tennenholtz demonstrate that it may be possible to prove whether that law is minimal and/or simple.

2.2 Agent Negotiation

Assuming an agent based system is social in nature, interaction will occur between agents in the system. Negotiation is one form in which this interaction between agents may take. Soh and Tsatsoulis (2001a) describe negotiation as the information exchange between two agents where an initiating agent wishes to persuade a

responding agent to accept a task. Thus, the goal of negotiation is to acquire a commitment at some level from another agent. Robinson and Volkov (1998) describe the negotiation process as being analogous to the software engineering life cycle in that agents involved in negotiation must analyze the problem presented to them, design a solution to that problem, and then implement the solution in some manner.

Perhaps one of the most widely investigated forms of negotiation is the contract net protocol, introduced by Smith (1980). In the contract net protocol, Smith states that negotiation has four important components: 1) it is a local process without centralized control, 2) information exchange is bilateral, 3) each party in the negotiation evaluates information from its own perspective, and 4) final agreement is reached mutually. The collection of nodes, or agents, in the system is termed the contract net. At any given time, an agent may take on the role of either the manager, who is requesting bids to perform a task, or a contractor, who submits a bid in response to a manager's request. The manager evaluates all bids received and chooses the most appropriate agent to perform the task (and receive the payment). The contract net protocol has been most widely used in systems of self-interested agents where individual agents have a mechanism for evaluating what actions may be in their own best interest in terms of their own goals and payoffs. However, at least in concept, it need not be limited to a self-interested system. Andersson and Sandholm (1998) have investigated one form of the contract net protocol, the leveled commitment contract in both self-interested and cooperative agents. In this case, the evaluation function a

cooperative agent uses to determine whether and what to bid is based on its own estimate of benefit to the overall system rather than just itself.

A generalization of the contract net protocol was introduced by Conry, Meyer and Lesser (1988), which they call multistage negotiation. In this strategy, agents iteratively exchange information such that overall solutions to the global problem are incrementally constructed. Responding agents propose solutions to a bid request and during the iteration process, they receive feedback concerning the impact their proposed approach will have on the overall system. This iterative exchange avoids the problems of subgoal interaction between partial solutions, determining whether a problem is overconstrained, and determining when an acceptable solution is found. As in the work of Andersson and Sandholm (1988), this approach was successfully applied to a cooperative agent society, although the authors state that it could be used in a self-interested society as long as agents were relatively honest.

Mouaddib (1997) also discusses an extension to the contract net protocol which he terms progressive negotiation, in which goals must be achieved by a given deadline. While similar to multistage negotiation (Conry, Meyer and Lesser, 1988), in its approach that agents may send back information on conflicts to a bidding agent, progressive negotiation decomposes problem structure hierarchically, and thus implements a form of partial goal planning (PGP, Durfee and Lesser, 1991), in a time-constrained domain. In order to meet deadlines, constraints on the problem are

ordered by a “preference criterion” which allows agents to address harder constraints first and provide a solution in a more timely manner.

The contract net protocol is certainly not the only approach available in negotiation. Faratin, Sierra and Jennings (1998) define a formal model of negotiation in which they define strategies and tactics that agents may take in generating offers, evaluating proposals, and offering counter proposals. A “negotiation thread” in this approach is the sequence of offers and counter offers in a two-party negotiation session.

Negotiation tactics is the set of functions that determines how to compute the value of a negotiation issue using a single criterion. The criteria are time, resources, previous offers and counter offers, and they play a role in time-dependent, resource-dependent and behavior-dependent tactics respectively. A negotiation strategy is a linear combination of tactics.

Lander and Lesser (1993) investigate negotiation in a heterogeneous agent environment where agents are cooperative. In a negotiation session, agents incrementally relax requirements for acceptable solutions, giving it similarity to multistage negotiation (Conry, Meyer and Lesser, 1988) and progressive negotiation (Mouaddib, 1997), however, it is not based on the contract net protocol. The negotiation model attempts to minimize the amount of information necessarily communicated between agents during a session. The negotiation operators used in their model include “initiate solution”, “critique solution”, “extend solution”, “relax

solution requirement” and “terminate”. They see this approach as one that provides enough generality for different application areas, and one that provides reasonable solutions in a reasonable amount of time.

A negotiation model based on argumentation was presented by Parsons and Jennings (1996) and Parsons, Sierra and Jennings (1998). Classical argumentation is a sequence of information leading to a conclusion. In this negotiation model, the initiating agent constructs an argument to pass on to another in a negotiation session. The receiving agent evaluates the argument and then either agrees with its validity, entirely disagrees, partially disagrees, or agrees but committing will cause the agent to be unable to achieve one of its own objectives. If an agent agrees, the negotiation is concluded positively. If a responding agent entirely disagrees, it may make a completely new suggestion by constructing its own argument to show why it disagrees. If there is only partial disagreement, the responding agent can send back an “undercutting” argument that applies to the point of disagreement only and the initiating agent can attempt to find an alternative. Finally, if committing to an act will cause an agent to miss its own objectives, replies with its own arguments for the objective and the initiating agent has the opportunity to propose something else that will satisfy both agents. The approach was described in a BDI agent architecture (Rao and Georgeff, 1991, 1995) and demonstrated how the two concepts could be integrated. Since the protocol communicates not only a request, but the reasoning

behind that request, it may be particularly helpful in heterogeneous agent societies, but may cause unnecessary overhead in homogeneous agent systems.

Kraus, Sycara and Evenchik (1998) also discuss negotiation through argumentation, and present a logical model and implementation approach. As before, their model is built upon the BDI agent model (Rao and Georgeff, 1991, 1995) where they define goals as a consistent subset of an agent's desires, and they add the concept of local preferences to an agent's mental states. Agents assign differing levels of importance to goals, and these goals motivate the agent's planning process. The planning process itself generates intentions which may either be "intend-to-do" (task level) or "intend-that" (goal level). Altogether, they term this agent model a BDIG model. The argument types that agents may use are: 1) threats, 2) future rewards, 3) appeal to a past reward, 4) appeal to precedents as a counter-example, 5) appeal to "prevailing practice", and 6) appeal to self-interest. These arguments are intended to influence another agent's beliefs and behavior. Although the work was set in a self-interested agent population, it's results may also be applicable, to a certain extent, in cooperative agent societies.

The unified negotiation protocol (Zlotkin and Rosenschein, 1991) is a protocol designed for negotiation between autonomous and noncooperative agents. They identify three types of interactions that can arise in a negotiation setting: 1) cooperation, 2) compromise, and 3) conflict. In a cooperative interaction, the

proposed request is preferred over individual goal attainment. In a compromise situation, an agent will agree to negotiate, though it may prefer to pursue goals on its own. However, the negotiation does present deals that are more beneficial than the agent acting alone. Finally, in a conflict situation, there are no deals in the negotiation that appear preferable to pursuing goals alone. The value of a deal is determined by “utility” which is defined as the difference in worth between the final state of the world to an agent and the cost that an agent must spend to bring about that world. Even in a conflict situation, the authors state that it is possible to reach partial cooperation on some parts of a negotiation.

In later work, Rosenschein and Zlotkin (1994) look at different negotiation protocols and evaluate the desirable attributes of each. One way they look at measuring the efficiency of a system is that of Pareto optimality, in which an agreement reached by agents can be made no better for any one of them without making it worse for one of the other agents. The two attributes they look at from this perspective are stability, that is, an agent has no incentive to deviate from a strategy, as in the game theory notion of equilibrium, and simplicity, which implies low computation and communication costs. They also prefer a distributed approach so that there is no central decision maker, and a symmetric system so that no one agent must play a special role. In a “best bid wins” scenario, the result is not stable, simple or efficient because bidding agents must expend resources to strategize about other agents’ behavior, and furthermore, the strategy may be incorrect, such that an agent who

could have provided the best solution overbids and loses the contract. In the “best bid wins, gets second price” approach, also called “Vickrey’s mechanism, the winning agent (the one with the lowest bid) gets paid the amount of the second lowest bid. This has the effect of separating who gets the bid from how much they get paid. In this approach, there is no incentive to either underbid or overbid, and it meets the criteria of stable, simple and efficient. The authors distinguish between three different levels of domains: task-oriented, state-oriented, and worth-oriented. In the task-oriented domain, agents have non-conflicting tasks to carry out, and these tasks can be redistributed to other agents, perhaps for reasons of overall system efficiency. The state-oriented domain is one in which goals represent acceptable final states. However, actions that lead to the attainment of goals may have side effects that may interfere with other agents’ tasks and goals. In the worth-oriented domain, goals specify final states, but each state has a rating as to its acceptability, and agents may not be able to reach their original goal (final state), but may still reach a state that is preferable to the starting state. In this work, utility is still measured as the difference between what would be gained by agreeing to a deal minus the cost or the total work performed. In going through these definitions, the authors stress the importance of evaluating different negotiation encounters with the domain type in mind such that the resulting strategy exhibits stability, simplicity and efficiency.

Soh and Tsatsoulis (2001a) describe a negotiation strategy that integrates case based reasoning with an argumentation based negotiation approach. In this work, the

underlying agent model is the BDI architecture (Rao and Georgeff, 1991, 1995). Agents in this system are autonomous, rational, communicative, reflective / aware, adaptive, and cooperative. Since the agent system is in a real-time environment, negotiations are constrained to be completed quickly. Since agents may play different roles, each agent has two case bases available to it, a negotiation initiating case base, and a responding case base. Case bases must be small, yet provide good coverage of possible cases, in order that negotiations may be completed in a constrained amount of time. Each case has three parts: 1) the problem space, or the beliefs of the agent at the time of the negotiation, 2) the solution space, or the strategy used in that session, and 3) the outcome of the session. Retrieval of similar cases is done based on a weighted similarity measure between current agent beliefs and the problem space of the case. The best case chosen is the one with the greatest degree of similarity, and if there are two such cases, the one chosen is the one with the better outcome. The best retrieved case is then adapted for the current negotiation problem and the strategy is then applied. Once a negotiation session is completed, the resulting case is added to the appropriate case base, either initiating or responding, dependent on the role of the agent in the negotiation. In this way, agents adapt over time to utilize better negotiation strategies.

2.3 Coalition Formation

Soh and Tsatsoulis (2001b) describe a coalition as a group of agents that collaborate to perform a task. A dynamic coalition is one that is formed to perform such a task and then dissolved once the task is complete. Coalition formation is necessary when any one agent lacks information, knowledge, or ability to individually achieve a goal

Coalition formation has been much studied in the context of game theory. Shenoy (1978, 1979) describes the theory of an n-person cooperative game as a mathematical theory of coalition behavior and then applies the mathematical analysis to the descriptive sociological theories of Caplow (1956) and Gamson (1961). Caplow's assumptions about coalition formation, in the context of a triad, are that members have differing amounts of strength and stronger members will try to control weaker ones, each member will seek to exert control over others, strength is additive and the strength of a coalition is equal to the sum of the strength of its members, and finally, if a stronger member attempts to coerce a weaker member into a non-beneficial coalition, it will force the formation of a coalition to oppose the coercion. Gamson's assumptions are that a member will prefer the cheapest winning coalition, the relationship in a two-player coalition must be reciprocal between the members, the process of coalition formation is a step-by-step process, and a two-person coalition essentially becomes the equivalent of one player with the strength of the sum of its members. Interestingly, Shenoy's (1978) game theoretic mathematical formulation

predicted the same outcomes as those predicted in the intuitive formulations of Caplow (1956) and Gamson (1961).

In more recent work targeted at coalition formation in agent systems, Tohme and Sandholm (1999) investigated the process with belief revision in bounded-rational self-interested agents. In their approach, agents make the decision to enter into a coalition dependent on the advantages they can gain from such a coalition. An agent's beliefs, at any given stage of the coalition formation process is a conditional probability distribution on the outcome, given previous steps of the process, and a coalition obtains stability when all of the member agents' beliefs converge. This is at the price of tractability, however; an optimal coalition that supports a Pareto optimal outcome can be exponential in the number of agents and the number of negotiation steps. In bounded-rational agents, however, the length of negotiation must necessarily be limited. If agents' beliefs do converge under the bounded condition, then stability has been reached in the coalition, but it may not converge if another coalition can block it in the same amount of bounded time. Given a greedy stepwise maximization of expected payoff, with deliberation and communication costs included, the authors demonstrated that stability can still be achieved in a finite number of steps.

Sandholm and Lesser (1995) also discuss coalition formation in the context of bounded rational agents. In the negotiation process, agents seek to minimize both the solution cost and the computation cost. They also use a "design-to-time" approach, as

opposed to an “anytime” framework (where the process can be stopped at any time and still provide a better result than where it was started), in that agents assume that the resulting solution of the coalition must fit within a certain time frame. In contrast to their later work (Sandholm and Lesser, 1996) they assume that coalition value may have a dependency on the actions of nonmember agents, and thus they model it as a normal form game, rather than as a characteristic function game. In varying the cost of computation, they found, not surprisingly, that higher computation costs led to smaller coalitions and vice versa. In a similar context of bounded rational agents, Larson and Sandholm (1999) looked at using three anytime structure generation approaches in coalition formation and found that none of the three algorithms (CSS1, MERGE, and SPLIT) emerged as being “best” between different problem domain types, however, the CSS1 algorithm did provide the most uniform performance across problems. They suggest that the design-to-time approach mentioned previously warrants further study.

In similar work, Sandholm et al. (1999) investigated the problem of coalition structure generation with worst case guarantees. They approach this by viewing all potential coalition combinations, in which coalitions must be disjoint, as a coalition structure graph, with one coalition of n agents at the bottom and n coalitions of 1 agent being at the top. First they search the bottom two levels of the graph to find “best” coalition groupings, and if time is left, they add a breadth-first (“anytime”)

search from the top of the graph. This approach rapidly decreases the worst-case bound found in the first part of the search.

Shehory and Kraus (1998) look at coalition formation under conditions where coalitions must be disjoint and also where they may overlap. They assume cooperative agents that attempt to maximize the common utility, as opposed to much of the game theoretic literature that assumes an adversarial role between coalitions. The agents in this approach are group-rational, meaning they will only form a coalition if the coalition benefit is at least as much as the sum of individual benefits would be. They are also bounded, in that coalition formation has an associated cost. Their algorithm is an anytime, distributed, greedy approach, and they argue that the anytime property is important since negotiations may have to be halted at some previously undetermined time. The main advantage to their approach seems to lie in the distributed nature of calculations in the system such that computation is more evenly distributed and unnecessary computations are avoided. While an optimal result is not achieved (this would be NP complete), it was shown that close to optimal results can be obtained.

The formation of coalitions by cooperative agents with incomplete information was investigated by Soh and Tsatsoulis (2001b). The dynamic and real-time characteristics of the environment further constrain the coalition formation process such that an optimal solution is impossible. The model used was an anytime

algorithm in which agents form a coalition “hastily” and then refine that coalition if time permits. The approach has three steps: 1) an initiating agent compiles a ranked list of other agents it believes would be helpful, 2) the agent begins a negotiation process with potential coalition members based on a case-based argumentation negotiation model (Soh and Tsatsoulis, 2001a), and 3) the initiating agent refines the coalition if it doesn’t satisfy its original goals and there is still time remaining to do so. An initiating agent forms a coalition only with its neighbors, a subset of the entire agent population about which the agent has some knowledge and with whom it can communicate directly. Some of the information about neighbors that an agent uses to evaluate them as potential coalition members include information about past encounters with the potential member, the current relationship to the member, and the initiating agent’s assessment of the potential member’s ability to perform the required task. The potential utility of a coalition is measured as a sum of all its members potential utilities plus the calculation of values from heuristics about the coalition size, flexibility or inflexibility of demands, and the priority of the event which must be dealt with. An initiating agent is bounded by the number of negotiations it can handle simultaneously, so the number of potential coalition members, at least on the first pass is also bounded. A coalition formed may be sub-optimal in that it is insufficient to meet all the goals of the initiating agent, but since members of the coalition have committed to the goals of the initiating agent, they will likely form their own coalitions and thus a sort of spreading activation will occur. Agents are not constrained to be members of only one coalition at a time. The overall effect of this

approach is to produce a “good-enough soon-enough” solution in response to a real-time event.

2.4 Commitment and Decommittment

The concepts of commitment and decommitment are highly intertwined with agent negotiation and coalition formation. In both negotiation and coalition formation, agents communicate with one another in order to get commitments, either to performing a task, adopting a goal, or agreeing to become a member of a coalition. Wherever a commitment exists, the potential for decommitment must also be considered, whether that decommitment is voluntary or forced.

The BDI agent work discussed previously (Rao and Georgeff, 1991, 1995) gives a basis for the concept of decommitment. In their discussion of intention, they describe the notion of a commitment to a previous decision. Commitment is what gives the system balance between reactivity and goal-directed behavior in an agent. In their view, a commitment has two parts, the “commitment condition”, which is the condition the agent has committed to maintain, and the “termination condition” which is the condition under which the agent may give up the condition. Since an agent does not have direct control over either its beliefs or desires, commitment can only apply to its intentions. Rao and Georgeff describe three types of agent behavior dependent on what types of termination conditions are allowed. “Blindly-committed” agents will

deny any changes to beliefs or desires that conflict with its commitments. A “single-minded” agent will allow changes to its beliefs and drop commitments accordingly. Last, an “open-minded” agent will allow changes in both beliefs and desires that may force commitments to be dropped.

It was also previously discussed that Schut and Woolridge (2001) extend the BDI model to include a decision theoretic approach to what they call “intention reconsideration”. Their approach allows an agent to decide for itself (as opposed to having an a priori commitment level) whether to perform an action or to deliberate over its intentions. Two conditions are defined under which reconsideration may be appropriate: 1) when conditions have changed such that an intention can no longer be achieved, it may be good to drop that intention, and 2) when opportunities arise it may be desirable to adopt additional intentions. While this approach moves the decommitment decisions from design-time to runtime, it only allows an agent to consider what action should be done “next”, a default action or a reconsideration. The approach also studied only the effects on a single-agent system, and thus still disregards effects on the system as a whole in a multi-agent system.

Again, as previously discussed, Norman, Sierra and Jennings (1998) defined a formalization of agent commitment. They discuss the fact that their approach does not account for decommitment, or “loss of commitment”, which they say may happen

because the action is no longer possible, the agents have agreed to be no longer bound to that agreement, or because the agreement has expired.

Meyer, van der Hoek and van Linder (1999) have also defined a formalization of agent commitment, but their model allows for the process of decommitment, at least under special conditions. In their formalization, an agent may decommit only if it no longer knows that commitment to be correct and feasible. This is still a rather restrictive condition, particularly in a time-bounded domain where events may occur that indicate an agent should perform one task over another previously committed task, even though the committed task is still correct and feasible. Since task performance is time dependent, the opportunity to perform the permitted task may expire. If the agent knows this may happen in advance, it would appear to be more beneficial to the system if it can inform any affected parties of the decommitment so that they may respond if time permits.

Panzarasa, Jennings and Norman (2001) address this aspect of decommitment. As one of their desiderata on social agent decision making they state that agents may fail in their cooperation. This is not expressed in terms of a forced failure, but in terms of agents not necessarily being benevolent or when an unpredictable event makes the commitment infeasible. In the authors' view, social and joint commitments are less strong than individual commitment or intention, unless additional conditions are imposed on joint intention. While they express this in terms of non-benevolent agent

societies, it could still occur in cooperative and benevolent societies since a cooperative agent has a goal of maximizing the overall system benefit, and if this creates a conflict with other committed goals or actions, decommitment will be the result.

Excelente-Toledo, Bourne and Jennings (2001) introduced the concept of differing levels of commitment, along with differing types of sanctions for decommitment, in a population of rational and self-interested agents. Experimentally, they showed that an intermediate degree of commitment, which implies a certain degree of loyalty, leads to better performance than either total commitment or loose commitment. Setting imposed penalties based on the prevailing context at the time of decommitment provided better performance than either fixed or partially sanctioned penalties. Because the agent population was self-interested, the goal of each agent was to maximize its own gain, however, which implies the concept may not necessarily directly extend to a fully cooperative agent population, in which the goal of each individual agent is to further the global goal of the entire system.

As previously discussed, Sen and Durfee (1996) describe the use of decommitment, or cancellation in their terms, to resolve conflicts in a meeting scheduling domain. In their discussion, cancellation occurs when another meeting of higher priority can only be scheduled at a time when a lower priority meeting has already been scheduled. In this approach, the cancellation of a single member of the meeting would cause the

cancellation and subsequent rescheduling of that meeting. Resolving conflicts through cancellation and rescheduling was left as an initial investigation in which they discuss some of the issues surrounding what they term a “significant problem”, and was not experimentally evaluated or validated. While they discuss what they term “commitment” strategies (Sen and Durfee, 1994, 1996, 1998), this is intended as a strategy to use in the scheduling process. One approach is that an agent blocks off and commits time to a tentatively scheduled meeting so that time is not available for other meetings, even though the scheduling has not yet been resolved. The other approach is that tentative, but as yet unscheduled, meetings are not committed until completely agreed upon. Although their terminology and problem characteristics differ from the focus of this research, their work provides some interesting parallels.

Chapter 3 – Formalization of the Problem

3 Approach

This research investigates the effects of both unilateral and negotiated decommitment in a collaborative agent society, under varying degrees of constraints. The theoretical framework underlying this research consists of several aspects of agent characteristics and behavior. First, a formalization of the problem of task scheduling is presented. Next, the basis for an individual agent is defined, and agents as a society are defined. Finally, agent interaction, including communication and negotiation are discussed, and the framework for both commitment and decommitment is defined.

3.1 Distributed Task Scheduling

An agent schedules tasks to be performed at a given point in time. These tasks may originate from the agent itself, or may be committed to at the request of another agent. The distributed task scheduling problem is defined by Sen and Durfee (1996) as a group of tasks to be scheduled requiring the use of a set of resources controlled by a set of agents. In their notation, the distributed scheduling problem is defined as $S = (A, T)$ where $A = \{a_1, a_2, \dots, a_k\}$, the set of agents with control of resources, and

$T = \{ \tau_1, \tau_2, \dots, \tau_n \}$, the set of tasks which may be scheduled. A task is represented by the tuple:

$$\tau_i = (A_i, h_i, l_i, w_i, S_i, a_i, d_i, T_i)$$

where:

$A_i \subseteq A$, is a set of agents controlling required resources;

$h_i \in A_i$, is the agent requesting the performance of the task;

l_i is the requested duration of the task;

w_i is the priority assigned to the task;

S_i is the set of possible starting times for the task;

a_i is the timestamp at which h_i requested the task be performed;

d_i is the deadline by which time the task must be scheduled;

T_i is the time at which the task is actually scheduled.

Although tasks may be defined a priori, this formalization does not require a static scheduling environment, and in fact, it is assumed that tasks arrive over time. In the multi-sensor target tracking domain, S_i represents the time at which a target is projected to first be visible to A_i , and l_i represents the duration of time it is expected that the target remains visible.

In this research, the task priority, w_i , represents an agent's local assessment of the estimated global utility of performing the task, or the potential *value* of making that commitment. This commitment value is defined as the tuple:

$$w_i = (p_i, v_{hi}, c_i, w_{hi}, dt_i)$$

where:

p_i is the default priority of that particular type of task;

v_{hi} is A_i 's assessment of the validity of h_i 's information;

c_i is the constrainedness of the task, comprised of the number of other agents also asked to perform the task and the duration (l_i) of the task;

w_{hi} is h_i 's assessment of the value of the task;

dt_i is the difference between the time the request was made and the requested start time, or $(S_i - a_i)$.

The first four factors, p_i , v_{hi} , c_i , w_{hi} , are applicable to evaluating task utility in multi-agent task scheduling, in general. The last factor, the difference between the time the request was made and the start time of the task, dt_i , is applicable to task utility in time-bounded domains.

3.2 Individual Agents and Agent Society

The individual agent model for this research is based on the BDI framework as described in Rao and Georgeff (1991, 1995), Parsons et al. (1998), and Noriega and Sierra (1996). As discussed previously, BDI stands for belief, desire and intention.

Beliefs represent the information an agent has about itself and the environment, including other agents in the system. Desires represent the motivations of the agent, and intentions express the deliberateness, or goal-driven nature, of the system. It is not required that desires be consistent, that is, an agent can desire one thing and also desire the opposite. Goals that an agent adopts based on desires and beliefs, however, must be consistent. Intentions represent the commitments an agent makes.

The agent characteristics that serve as a basis for this research are that agents are:

1. Collaborative and Benevolent – An agent attempts to act and plan according to the goal of maximizing the benefit of the entire system, rather than take the self-interested position of maximizing its own personal profit. Since agents are benevolent, honesty follows as a characteristic. They do not misrepresent their beliefs, desires or intentions in their interactions with other agents. It is not required, however, that agents are always correct. They may have false beliefs that they believe are true, and these may be communicated to others.
2. Rational – An agent does not act or plan in a manner that is detrimental to the agent society as a whole, according to its own local estimate of the global utility of a plan or an action.
3. Autonomous – An agent may take the arguments of another agent under consideration in making a decision, but ultimately, it is the agent itself that decides whether to perform an action or not.

4. Communicative – Agents have the ability to communicate information to others and to receive information from others.
5. Multitasking – Agents are capable of performing certain actions in parallel. This allows an agent to reason about future actions while performing a current action, and to communicate with others concurrently.
6. Capable of Negotiation – An agent may initiate negotiations with another agent, and an agent may also respond to another agent's request. Negotiations may be either a request or a communication of the intention to drop a commitment that affects another agent.
7. Capable of Time-Dependent Response and Planning – Agents exist in a time-bounded environment in which they must react to events that occur at the present moment and plan for events that are expected to happen at some point in the future. In this sense, they are both reactive and proactive.
8. Capable of Learning – Over time, agents build up information about other agents and the environment and may modify their actions based on past information.

3.2.1 Agent Society and Interaction

Following the research of Soh and Tsatsoulis (2001b), the social agent environment can be expressed as follows:

Ω – a multi-agent system

Ψ – a “neighborhood” in the system

$\lambda(\alpha, \beta)$ – predicate indicating agent α knows about agent β

A neighborhood is defined as:

$$\Psi \subseteq \Omega, \Psi \neq \emptyset$$

Every agent, α_i , in neighborhood Ψ , knows about all other agents in that neighborhood. That is:

$$\lambda(\alpha_i, \alpha_j) \forall i \forall j \alpha_i, \alpha_j \in \Psi$$

There may be any number of neighborhoods, they need not have the same number of members, and neighborhoods may overlap. Thus,

$$\Omega = \{\Psi_1, \Psi_2, \dots, \Psi_N\}$$

3.3 Negotiation

Agents may communicate with other agents within their neighborhood, that is, an agent may only communicate with those that it is aware of. Negotiation requires communication, thus, negotiation is also restricted to agents within the same neighborhood.

In initiating a negotiation session, an agent may either be requesting that another agent perform a task, or the initiating agent may be informing another agent that it is contemplating dropping a commitment that concerns the other agent, and is asking for

feedback on the other agent's perception of the importance of that particular commitment. When an agent initiates a negotiation as a request, it is the responding agent that holds the power over whether that request will be honored or not, therefore the initiating agent must provide convincing arguments in favor of the request. On the other hand, when an agent initiates a negotiation about dropping a previous commitment, it is the initiating agent that makes the final decision over whether the commitment will be dropped or not, and the responding agent must provide arguments to convince the initiator otherwise, if the responding agent still wishes to have its original request honored.

Since agents are not omniscient, each agent has its own local estimate of the global utility of a given action or plan. It is this estimate that is used to request cooperation from another agent, who in turn will evaluate the utility of that request from its own perspective in determining whether to accept the request or not. Conversely, an agent will use the utility estimate to determine if a commitment should be dropped, and if so, communicate this to affected agents. Agents responding to a decommitment negotiation session will use their own utility estimates as return arguments.

An agent stores information about other agents, and it may use this information to determine candidates that are more likely to agree to a request, and thus determine which agent(s) to negotiate with. Over time an agent will build an information base about its neighbors, including: 1) how often a neighbor agreed to a request, 2) how

often that neighbor honored that request, 3) how often the agent itself agreed to a neighbor's request, and 4) how often performing a task for a given neighbor resulted in performing a valuable service, in this case, detecting a target. In this manner, an agent learns which of its neighbors are more reliable and which neighbors tend to provide valid information. The agent will base future negotiation considerations on this information.

One of the system characteristics of interest is time-bounded response and planning, thus negotiations are time-limited. A negotiation session may be terminated if it is sensed that there is not enough time to complete it successfully. The amount of time available to a session is dependent on the expected future time the some action must occur, along with an agent's assessment of the likelihood of success at that point in the negotiation.

3.3.1 Agent Commitment and Decommittment

An agent may have its own commitments that do not affect others. Dropping a commitment that doesn't involve others is simply a matter of determining whether another task has a higher estimated utility. Agents also take on commitments through negotiations with other agents, that is by agreeing to requests.

The *value* of a commitment is determined by an agent's local estimate of that commitment's global utility, as previously described. The *strength* of a commitment is a measurement of the impact it has on other agents. When decommitment is done as a rational act, that is, the agent makes a decision to drop a commitment, that decision is based on the value and the strength of the commitment.

Should an opportunity arise where the agent may perform an action that it perceives as more beneficial to the social welfare than what it had planned (for example, it receives a request to perform a more highly valued action), the agent may tentatively decide to drop the previous commitment. This may be negotiated with any agents impacted by the potential decommitment, and the agent may choose not to decommit, but to decline the new opportunity instead. A similar situation occurs in which an agent may perceive that a scheduled task is no longer necessary. The agent may then negotiate to drop that commitment with those affected.

The strength of a commitment is defined as a tuple:

$$\text{str}_i = (n_i, r_{hi}, \text{dnow}_i)$$

where:

n_i is the number of agents potentially affected by the decommitment;

r_{hi} is the perceived reliability of the neighbor to whom the commitment was made, that is, the number of times that neighbor honored commitments to A_i ;

dn_{ow_i} is the difference between the scheduled start time of the task and the current time.

Using negotiation as part of the decision process to decommit or not allows the decommitting agent to make that decision based on more information than what it may have locally. It also allows the affected agent to re-evaluate the value of its previous request and to provide arguments for keeping that commitment if it determines that it is still important. Thus, a negotiated decommitment takes into account a somewhat more complete estimate of the task's utility, and provides a better basis for maximizing social welfare than unilateral decommitment would.

Chapter 4 – Implementation

4.1 Problem Domain

The Autonomous Negotiating Teams (ANTS) project, a DARPA funded research effort led by Drs. Tsatsoulis, Niehaus and James, of the Information and Telecommunication Technology Center (ITTC) of the University of Kansas, is an agent-based system that provided the background programming environment for this research. The domain area in this research is that of multi-sensor target tracking in which each agent has one sensor (radar) with three sectors that allow it a 360° view of a distance-limited range. Only one sector may be active at any given time. As a target moves through the area, agents controlling sensors in different locations must cooperate to track the target as it moves in and out of their sensing ranges. Agent cooperation requires inter-agent communication. While the global problem is one of tracking a target through space, each agent has only a local view of the situation and therefore must negotiate with others to perform tasks it cannot perform alone. Furthermore, a moving target implies that the problem is time-constrained, so agents must negotiate and act within a reasonable amount of time, relative to the target speed.

Measurements taken by a sensor can only indicate that a target is present within its sensing area, which provides a general area within which the target is located.

Additional measurements are required, from other sensors, in order to get a better estimate of target position and velocity. Three measurements from different sensors within a two second time period are needed in order to provide good estimates, however two measurements will at least narrow the possibilities of position and velocity of the target. Agents must therefore communicate and negotiate with each other to coordinate measurements for the best estimate possible. This communication occurs for both the condition where a target has been detected now (reactively), and also for where the target is expected to be in the future (proactively), based on current target estimates.

Several programs are used in conjunction with the agent system. Two of these were developed outside the University of Kansas. The first is a radar simulator, `RadSim`, which provides simulated amplitude and frequency measurements of a target when a given sensor's sector is activated. `RadSim` was developed by the Rome Air Force Laboratory. The second is the `tracker`, developed by the University of South Carolina, which, given measurements obtained by sensors, will attempt to determine the location and velocity of the target. Both of these programs are written in the Java programming language.

Figure 4.1 shows a screen from the radar simulator, `Radsim`. In this picture, there are six sensors, depicted by the darker circles labeled S1 through S6, and two targets, depicted by the lighter circles labeled M2 and M3. The wedge shapes connected to each sensor show which sector is active, though the sensor cannot actually sense within that entire range. Each target is shown along with its track (the oval shapes in the screen shot), and its current position. The lower left hand corner shows a wedge shape emanating from “sensors” at position (0,0). These are not actual sensors, however, but represent the `tracker` assigned to each target. While running the program, the wedge shapes depicting sectors are shown in yellow when an agent is calibrating that sector, and in blue when an agent is taking amplitude and frequency measurements.

Programs written at the University of Kansas under the ANTS project, and which are used in the current research, are the `proxy`, the `tracker agent` program, and the `radar model`. The `proxy` program provides socket communication capability between the Java code in `Radsim` and the `tracker` and the C++ agent code. The `tracker agent` provides an interface to the `tracker(s)` used in the program, and is written in Java. Finally, the `radar model` program approximates the sensing area of each sensor’s sector with an ellipse, and given current estimates of a target’s location and velocity, determines which sensor/sector pairs can see the target at present, and which ones will see it at what point in the future, assuming that the input

is correct. Figure 4.2 shows the relationship of the agent software used in this project to the external programs.

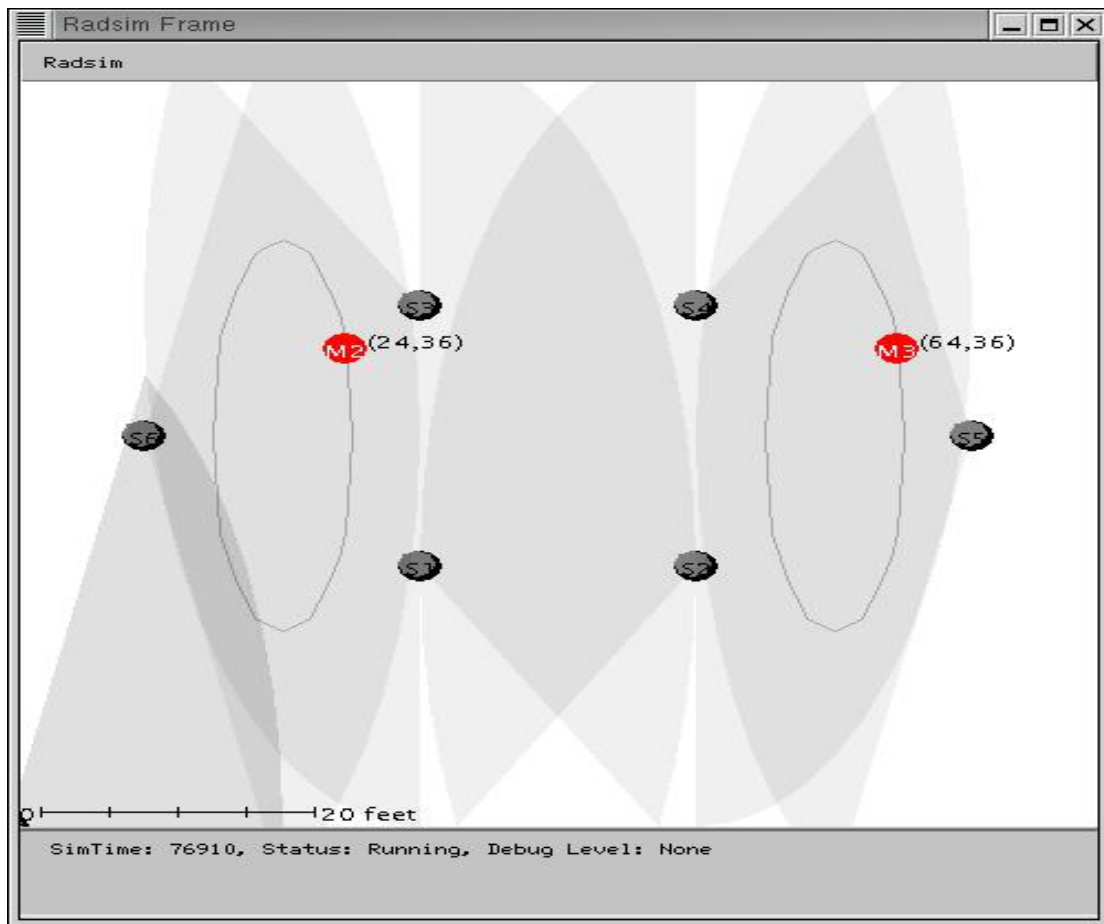


Figure 4.1: Radar Simulator Screen; 6 Sensors, 2 Targets

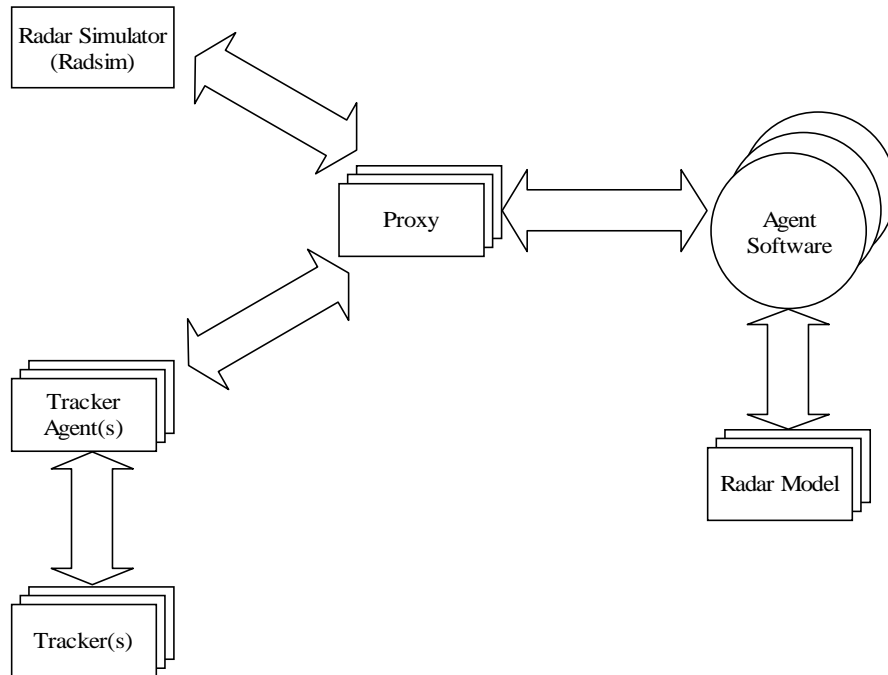


Figure 4.2: Agent Software and External Programs

When running a given experimental condition, one copy of the radar simulator executes, however, there is one tracker and associated tracker agent for each target in that particular experimental condition, and similarly one proxy and one radar model for each agent in the test condition. The entire configuration runs in a networked environment, in this case, using the Red Hat Linux Version 7.3 operating system.

4.2 Agent Software

The current agent software is designed to run in one of three experimental modes: 1) baseline, 2) unilateral decommitment, and 3) negotiated decommitment. In the baseline mode, agents negotiate to make commitments, but do not drop commitments. The unilateral decommitment mode includes the baseline processing, with the addition that if the agent finds another task of more utility than a previously made commitment, it will drop that commitment and notify the affected party. Finally, the negotiated decommitment mode is similar to the unilateral decommitment mode, however, an agent will not unilaterally drop a commitment, but will make that decision after negotiating with the affected agent. How decisions are made are dependent on the experimental mode.

4.2.1 Agent Characteristics

Agents are characterized by the following traits (Soh and Tsatsoulis, 2001a):

1. Autonomous – Agents run without human interaction and control their own actions without explicit direction from other agents.
2. Rational – Agents have goals and can reason about actions to achieve these goals. Furthermore, they do not behave in a manner inconsistent with goal achievement.
3. Communicative – Each agent may both initiate messages to and respond to messages from other agents, providing a basis for negotiation.

4. Reflective – An agent is aware of the current situation, including its observations, its information state and assumptions, and any communications from other agents, and makes decisions based on this awareness.
5. Adaptive – As the situation or environment changes, agents learn to perform better based on past experience.
6. Cooperative – Agents are biased to cooperate with other agents in the system in order to solve the global problem.
7. Homogeneous – In this system, all agents have the same capabilities and communication methods. Each agent may encounter differing local conditions, however, and as adaptation occurs over time, the knowledge of different agents may vary.
8. Time-Bounded – Agents operate in a time-bounded environment, performing actions both reactively and proactively.

4.2.2 Agent Architecture

As described in Soh and Tsatsoulis (2001a, 2001b) the agent architecture is based on the BDI framework as defined in Rao and Georgeff (1995), Parsons *et al.* (1998), and Noriega and Sierra (1996). This architecture was retained in the restructured agent software developed for this research. Agents must communicate with each other and with the radar simulator, and they must maintain a schedule of upcoming tasks. Since

events triggering the performance of these tasks is asynchronous, agents were implemented using a multi-threaded approach.

Upon execution, each agent spawns a thread for its negotiator, communicator, and the scheduler. The agent thread then acts as the main decision maker and informs its other threads when non-default actions must be performed. When a target is detected, the main agent thread decides whether to suspend currently planned activities to track the target, whether (and which) agents to inform that the target is currently within their range, and which agents to negotiate with in the future. Figure 4.3 shows this threading relationship diagrammatically. All threads execute until they are notified that the simulation is ending, at which point they clean up internal data, close any files and terminate gracefully. The main agent thread allows time for each of its spawned threads to shut down and then terminates its own execution.

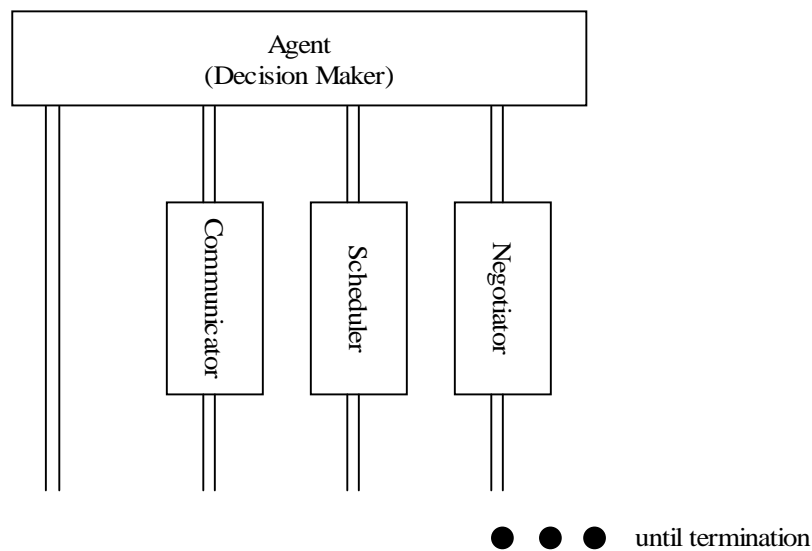


Figure 4.3: Agent Multithreading




Figure 4.4: Communicator Operation

The communicator thread is responsible for sending messages to and from the radar simulator, other agents, and the tracker agents, via the proxy program. It removes messages from an outgoing mailbox and directs them to the appropriate entities. The communicator also receives messages from external entities, formats them into the internal agent message structure, and places them in the incoming mailbox so that the appropriate thread may pick up the message and process it.




Figure 4.5: Scheduler Operation

The scheduler thread is responsible for maintaining the schedule. It monitors the current simulation time and scheduled sensor tasks and places current tasks in the outgoing mailbox for the communicator to retrieve and send to the radar simulator. The scheduler also adds default tasks to the schedule when the schedule spans less than the minimum time. The standard tasks that the scheduler works with are the idle task, calibration and detection. Calibration involves taking measurements at different sectors when a target is not present to determine the background noise level. Detection is the process of taking measurements on different sectors when there is a possibility that a target is present. The idle task is self-explanatory. A fourth type of task is tracking, in which a target has been detected, and the sensor continues to take measurements which are then sent to the tracker agent. This task is only scheduled when a target has been detected, however, and the decision to schedule this task rests with the decision maker (main agent thread).

The negotiator thread receives notification from the main agent thread when negotiations should be performed, that is, when a target has been detected and it is projected to pass within the field of vision of neighboring agent sensors within the near future. It formats negotiation requests to other agents into message format and places them in the outgoing mailbox for the communicator to send. It also retrieves messages from the incoming mailbox and processes these as either requests from other agents or responses to its own requests. A request is a communication issued by

an agent either to ask for assistance or to decommit from a previous commitment. A response is the return message from either type of request.




Figure 4.6: Negotiator Operation

The agent thread acts as the decision maker, and thus may access the scheduler's schedule directly in order to add or delete scheduled tasks and to make decisions concerning the value of requests from other agents. Note that the negotiator thread can record responses from other agents in the form of commitments from them, but only the decision maker / agent thread can decide to make a commitment to another agent, or to drop a previous commitment to another agent.

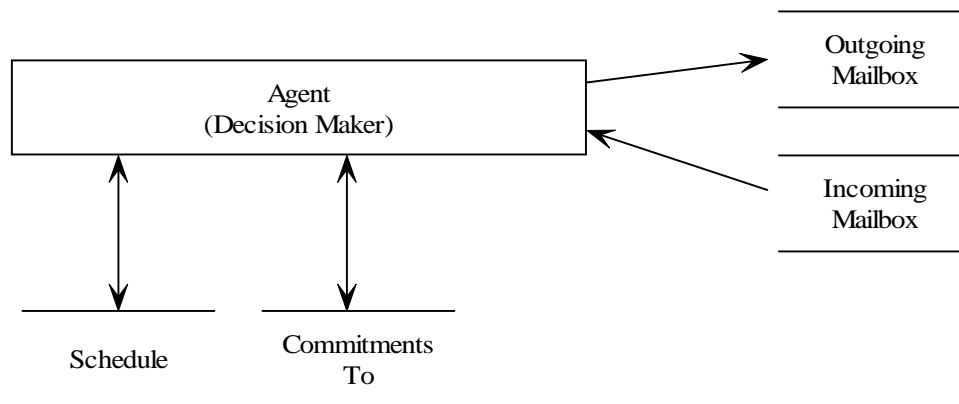


Figure 4.7: Decision Maker (Agent) Operation

4.2.3 Agent Interaction

As described in Soh and Tsatsoulis (2001a, 2001b), agents negotiate with each other to perform tasks in order to meet the global goal of tracking a target. During a negotiation session, agents exchange information (“beliefs” in the BDI model) and attempt to persuade others to commit to what the requestor is asking for. Because time is a bounded resource in this domain, it is important that negotiations are resolved quickly enough that a task may be carried out while it is still pertinent.

In the multi-sensor agents domain, neighborhoods are determined by sensors that are physically close to one another such that their sensing areas overlap. Agents are thus considered neighbors if they within a specified distance from one another. Agents that

negotiate with each other will do so only with those in their own neighborhood. In the domain problem, when an agent detects a target, it tracks the target locally to determine an estimate on the direction and velocity of the target's movement and projects where that target may be in the future. The agent then identifies other agents that it knows have sensor coverage of the projected path of the target and negotiates with those agents to perform tracking tasks when the target is projected to be in their range.

Since, by definition of agent characteristics, agents are known to be cooperative and benevolent, one can assume that if an agent can commit to performing a task at the requested time, for the requested duration, it will do so. If the agent is unable to commit to the full requested duration, but is available for a subset of that time, one can assume it will counter with the largest time span it is able to commit to. In this implementation, an agent sends all known information to a neighbor when initiating a negotiation. Under these circumstances, multi-stage negotiations in requests for assistance are unnecessary, and the negotiation strategy is to simply make requests containing the information that the initiating agent believes to be true.

The negotiation protocol used in the current agent system for assistance requests is shown in Figure 4.8. In this diagram, final states, or outcomes of negotiation are represented as squares. The initial state is represented by a double circle and the intermediate state by a single circle. A negotiation may ultimately succeed, fail, or

time out. Negotiation proceeds by either an immediate agreement or refusal of the request by the responding agent, or by a counter offer from the responding agent, representing a subset of the original time requested. These actions immediately result in either the success or failure of the negotiation. Another condition of negotiation failure is if the initiating agent does not receive a response within the time frame necessary to perform the task. In this case, the negotiation is considered to have timed out.




Figure 4.8: Request Negotiation Protocol. Squares represent final states and double circles represent initial states.

Unlike requests for assistance, negotiated decommitment requests may require more than one stage. When an agent requests to decommit from another agent, if the responding agent either agrees or doesn't respond within the allotted time frame, the initiating agent will drop that commitment. If, however, the responding agent responds with updated information, the initiating agent will re-evaluate its decision to decommit with respect to the new information. The initiating agent may agree with the responding agent, in which case it will retain its original commitment. Or, it may either deny the counter proposal or run out of time, which by default results in dropping that commitment, as shown in Figure 4.9. Unilateral, or non-negotiated decommitment reduces to simply notifying the affected agent of the intent to decommit.

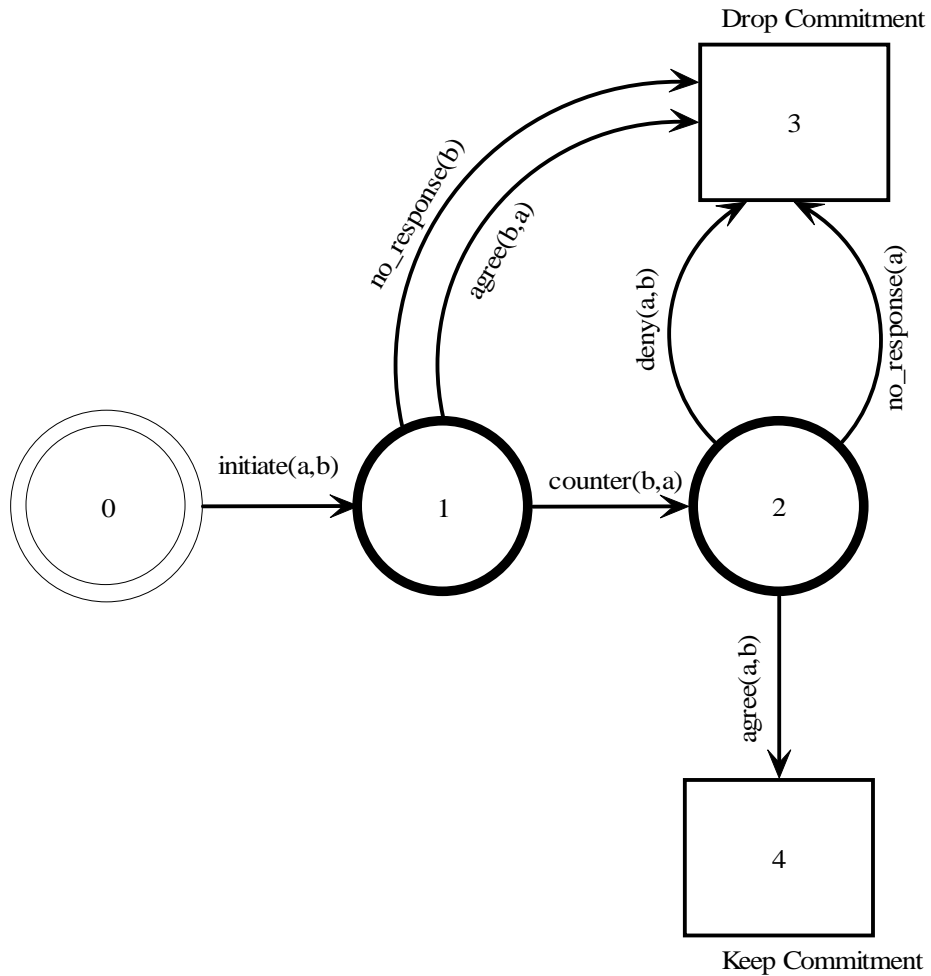


Figure 4.9: Negotiated Decommittment

4.2.4 Local Estimate of Global Utility

The utility of a task is evaluated by the decision maker, or main agent thread, to assess the potential value of performing that task, or if requested by another agent, the value of making that commitment. The commitment value is used to assess the utility

of a potential commitment while the commitment strength is used to modify the commitment value, or utility, when considering decommitment.

The commitment value, or estimated utility of performing a potential task, was previously defined as:

$$w_i = (p_i, v_{hi}, c_i, w_{hi}, dt_i)$$

where:

p_i is the default priority of that particular type of task;

v_{hi} is A_i 's assessment of the validity of the neighbor h_i 's information;

c_i is the constrainedness of the task, comprised of the number of other agents also asked to perform the task and the duration (l_i) of the task;

w_{hi} is h_i 's assessment of the value of the task;

dt_i is the difference between the time the request was made and the requested start time, or $(S_i - a_i)$.

In the domain of multi-agent target tracking, these components can be defined more concretely as:

1. Default task priority: The implicit goal of the agent system is to track targets, and this is reflected in the default priorities assigned to tasks, with tracking being the

highest, followed by search and detect, then calibration, and last, the idle task.

These are implemented as constant values, ranging between 1 and 10.

2. Confidence that a target will be in a given sector at a given time: Results returned by the `tracker` software are not necessarily accurate, and are particularly suspect toward the beginning of a tracking event since the tracker relies on past information to predict. These location and velocity estimates are what are used to project future target locations, though. In this implementation, an agent keeps track of each time it performs a tracking task at the request of another agent, and how often the performance of that task results in detecting a target. The result is an information validity rating for each neighbor which is used in assessing the utility of future requests from that same neighbor. The higher the validity rating of the neighbor, the higher the utility of the task.
3. Number of constraints on a task: Two conditions are taken into consideration when determining the constrainedness of a task. The first is the requested duration of the task, which is the estimated visibility window produced by the `radar` model. The smaller the visibility window, the more constrained the task, and the value of performing it is increased. The second condition is the number of other agents also asked to perform this same task. When there are fewer others to assist in tracking, the utility of taking on the commitment increases.
4. The requesting agent's utility assessment of the task: Before making a request, the requesting agent evaluates its own utility of the task based on the information it has

available. The responding agent uses the requesting agent's assessment as one factor in its utility assessment.

5. Nearness in time of a track task to the time it was requested: The closer in time that a track task is to be performed to the time of the request, the more likely that the tracking information is, if not correct, at least closer to actual values, than if it was projected far into the future. Thus, a smaller difference in time implies the utility of the task is higher.

When an agent is considering dropping a commitment, it reassess the commitment's utility value, and also assesses the commitment strength. The commitment strength may either increase or decrease the perceived utility of keeping that commitment. The commitment strength of a task was previously defined as:

$$\text{str}_i = (n_i, r_{hi}, \text{dnow}_i)$$

where:

n_i is the number of agents potentially affected by the decommitment;

r_{hi} is the perceived reliability of the neighbor to whom the commitment was made, that is, the number of times that neighbor honored commitments to A_i ;

dnow_i is the difference between the scheduled start time of the task and the current time.

The implementation of these factors in the target tracking domain is operationalized as:

1. Number of agents: One of the pieces of information transmitted in a negotiated request for assistance is the number of agents that share the same visibility window in tracking the target. The more agents that may be affected by a decommitment, the less attractive it is to decommit.
2. Relationship to affected agents: An agent keeps track of how many times a neighbor has agreed to a commitment and then later has dropped that commitment, and thus develops a reliability rating of each neighbor. An agent is more willing to drop a commitment with an agent that has proved unreliable in the past than it is to drop one with a reliable agent.
3. Nearness in time to “now”: The closer in time it is that a task is scheduled to be performed, the more costly it will be to drop that commitment because it will be more difficult for affected agents to compensate.

4.2.5 Decision Criteria

The decision maker, or main agent thread, is responsible for all commitment and decommitment decisions. These can be divided into decisions based on information it receives from its sensor, from a tracker agent, or from another agent. Information from the agent's sensor will either indicate a target has been detected or not.

Information from the tracker agent is in the form of updated location and velocity estimates at a given time. Finally, information coming from another agent may be a request for assistance, a response to a request, a request for decommitment, a response to a decommitment request, or unilateral notification of a decommitment. How the decision is made depends upon the current experimental condition being run: 1) baseline experiment, or negotiation for commitments with no decommitment allowed; 2) unilateral decommitment experiment, or negotiation for commitment, but decommitment performed without negotiating; or, 3) negotiated decommitment experiment, using negotiation for both committing to and decommitting from obligations.

4.2.5.1 Incoming Sensor Information: “Track Now” Task

When an agent receives information from its sensor that a target has been detected now, it must decide whether to track that target now or not, dependent on what task it is performing, and what tasks it is scheduled to perform in the near future. A potential “track now” task has the highest default priority, since the agent knows with certainty that a target is present. The agent requires a minimum of two seconds to dedicate to tracking a target, and would prefer to track it for five seconds.

In the baseline experimental condition, if the agent has at least two seconds of time on its schedule containing tasks of lower utility that were all originated as its own tasks,

it will decide to track the target now, taking up to five seconds if its schedule permits. However, if there are pending commitments on its schedule to other agents, whether these are of lower utility or not, it will honor those commitments, resulting in either a lesser duration for the “track now” task, or in simply not tracking the target now.

With a “track now” task, both the unilateral decommitment and the negotiated decommitment experimental conditions use the same decision making process. That is, the negotiated decommitment condition reduces to unilateral decommitment since there is no time to negotiate a decommitment. Recall from the negotiation protocol that a negotiated decommitment may reach a time out condition, in which case the agent unilaterally decommits. In a “track now” situation, the time out condition occurs immediately. For a “track now” task, then, the agent checks the schedule for any upcoming tasks that are of lower utility than the “track now” task, and performs the “track now” task if it has at least two seconds to do so. If there are commitments to other agents of lower utility on the schedule, the agent sends a unilateral decommitment notice to all those affected. Furthermore, the agent will also check what task is being performed now, and if it is of lower utility, it may send the communicator a message to stop the current task in favor of tracking the current target.

4.2.5.2 Incoming Tracker Information

When a tracking task is performed and a signal is detected, indicating that a target is present, that information is automatically sent to the `tracker agent`. Periodically, the `tracker agent` sends messages back to the agent to relay its current estimates of target location and velocity. When the agent thread receives the new estimates, it then sends this information to the `radar model` program, which projects the trajectory of the target in the future, as a straight line, and determines which of the agent's neighbors have current visibility of the target now, and which will have visibility later, in addition to what the projected visibility window is.

From the `radar model` information, the agent notifies those neighbors to which it believes a target is visible now, and requests that they track the target now. It also requests assistance from neighbors in the future, given the projected entry and exit times of the targets in a neighbors visibility field.

4.2.5.3 Incoming Agent Information

Information received from other agents includes: 1). requests to track now, 2). requests to assist later, 3). requests to decommit, 4). responses to requests to track now, track later or decommit, or 5). notification of unilateral decommitment. In general, responses from another agent, and notification of a unilateral decommitment,

require no decision-making, and are thus sent directly to the negotiator in order to record interactions with its neighbors. The exception to this is a response to a negotiated decommitment request, which will be covered in detail later. The information that the decision-making agent thread must deal with, then, are the requests to track now, track later or decommit.

Request to Track Now: A request from another agent to track a target now means that the other agent has received updated target location and velocity estimates from a tracker agent, and that its radar model program indicates the target is visible to me now. This condition is very similar to a “track now” task resulting from information obtained from my own sensor, however, the information validity is more suspect. The requesting agent, of course, believes the information to be true and accurate, but the receiving agent must re-evaluate the information based on the validity of the information it has received from the sending agent in the past, thus re-evaluating the utility of the potential commitment. Given a re-evaluation of its utility, it is treated exactly like the “track now” task discussed previously. That is, under the baseline condition, the agent will accept the task if upcoming scheduled tasks are of lower utility, there is enough time in the schedule, and the upcoming tasks are not commitments to other agents. Under the unilateral and negotiated decommitment conditions, if there are lower utility tasks scheduled for other agents, these commitments will be dropped and the other agents notified. The responses the agent may make are to agree, meaning that the full requested duration has been accepted,

counter offer, meaning that the full duration can't be committed to, but a subset of that time is available and has been committed, or refuse, meaning that the commitment cannot be made.

Request to Assist Later: A request from another agent to track a target at a later time is handled in the same manner as a “track now” request in the baseline and unilateral decommitment experiment conditions, looking at the part of the agent's schedule for which the task would be scheduled if it is committed to. The utility value of this request is expected to be lower, since the request concerns projected estimates of target location, but it is evaluated against scheduled commitments in the same manner. However, it is handled differently in the negotiated decommitment experiment condition.

Under the negotiated decommitment condition, the agent will look at its schedule for the time that the potential task would be scheduled, and if there are commitments to other agents with lower utility than the potential task, the agent will request a decommitment from each of the affected agents. Before making a response to the requesting agent, the responding agent will wait for responses to its decommitment request from all affected agents up to a point where the negotiation times out. If it times out, this reduces to a unilateral decommitment, all affected agents are notified of the unilateral decommitment, and the original requesting agent is notified of either an agree or counter offer response. (If this response would have been a refusal from

the beginning, the process of requesting decommitment would not be initiated.) If all affected agents respond before the timeout occurs, the agent must make further decisions, as described in the response to decommitment section.

Request to Decommit: If an agent receives a request to decommit from another agent, it re-evaluates the commitment that the requesting agent wishes to drop, using the latest information it has available from both the `tracker` agent and the `radar` model. If this results in a lower utility than what the commitment had previously, the agent agrees to the decommitment and notifies the requesting agent. If it results in a higher utility, the agent makes a counter offer to the requesting agent. In either case, the agent sends the re-evaluated utility to the requesting agent so that the requesting agent may modify the utility of that commitment on its schedule to reflect the latest information.

Response to Decommitment Request: When an agent requests a (negotiated) decommitment from another agent, it waits for a response until the negotiation times out, or until all agents have responded. If each affected agent agrees to the decommitment, this in effect, reduces to a unilateral decommitment in that all affected agents, who have agreed, are notified that the commitment is being dropped, and the agent that requested the commitment which started this process is notified of either agreement, if the full requested duration can be scheduled, or a counter offer, if only a subset of that time may be scheduled. This means that the commitment now

made under the negotiated decommitment condition has the same utility as if that commitment had been made under the unilateral decommitment condition. However, when the affected agents inform the agent of their agreement to the decommitment, they have re-evaluated the utility of the previous commitment and will only agree if it turns out to have lower utility. This has the effect of reducing the utility of the commitments that had been scheduled.

If any of the affected agents makes a counter offer to the request for decommitment, and this counter offer has a higher utility than the requested task the agent has been considering, the agent will accept the decommitment counter offer and refuse the requested task. If the counter offer has lower utility than that of the commitment under consideration, however, the agent will treat that much the same as an agreement to the decommitment, and once all affected agents have responded, it will unilaterally decommit, notify the affected agents of that decommitment, and notify the requesting agent that it has accepted its request.

Chapter 5 – Experimental Design

Based on the hypotheses of this research, experimental conditions were designed to measure the effect of decommitment. Recall that the hypotheses were:

Given a society of agents that are characterized as autonomous, rational, collaborative, benevolent, and capable of negotiation and planning time-dependent activities:

1. Allowing agent decommitment will improve the overall goal achievement of the system.
2. When an agent's decommitment to a task or goal affects other agents, negotiating the terms of the decommitment with those other agents will be more beneficial than unilateral decommitment, despite the associated increase in overhead. The improvement is expected as a result of a more refined estimation of global utility when more than one agent contributes to that estimate.
3. As the number of constraints on the system increase, the overall goal achievement of the system will degrade gracefully.

Thus, the measurement of interest is the overall goal achievement of the system under the differing conditions of negotiation without decommitment, unilateral decommitment, and negotiated decommitment.

5.1 Performance Evaluation Criteria

In the multi-sensor target tracking domain, the overall system goal is accuracy in tracking moving targets. There are several difficulties in using tracking accuracy as the only performance measure, however, since the tracking software is external to the agent software, and we have no control over its operation. The tracking software operates by estimating the location and velocity of a target by using three amplitude and frequency measurements from different sensors during a two second window. It has not been shown to provide overly accurate predictions of target location and velocity, particularly in early stages of a program run, or when a target changes direction (personal communications, Tsatsoulis, 2001). Furthermore, degradation of performance of the tracking software is not graceful. When predictions begin to fail, they begin to fail catastrophically. Finally, the tracker software is not a program developed locally, and therefore improvements to the software are not within the scope of this project. Because of these issues, overall tracking accuracy of the system was measured under all experimental conditions, but it was not used as the *primary* evaluation criteria for overall goal achievement.

Subgoals of the ANTS system are to provide the tracker software with timely measurements of detected targets simultaneously from different locations so that the tracker can do the best job it is capable of. These subgoals are under the control of the ANTS software, therefore these were used to evaluate the research hypotheses. In

particular, under each experimental condition, the following measurements were made:

1. Number of measurements planned per target: The more measurements made of a target, the greater the chance for better tracking. Planned measurements are defined as those that were made as a result of a decision by the agent, and not measurements resulting in target detection that were made during a random search and detect task.
2. Number of times three or more measurements were taken in a two second window per target: The `tracker` software requires three measurements within a two second window in order to perform triangulation and produce its best possible estimate of target location. These measurements must be taken from different sensors, implying successful agent cooperation.
3. Balanced measurement between multiple targets: When multiple targets are present, agents should take measurements of all targets approximately equally in order to track each target adequately. This measurement is only applicable in experimental conditions that have multiple targets.
4. Total number of target measurements made: This includes both planned and unplanned measurements made when a target was present. It should be noted that several sensor measurements may be taken for a single planned measurements. That is, the sensor will repeatedly take measurements for the duration of a tracking task at regular intervals. The more total measurements taken , the better the chance of producing reasonable tracking estimates.

5. Average utility per millisecond: For each experimental condition, during each run, the average utility of the existing schedule, the existing schedule as would be modified by feedback in a negotiated decommitment experiment, the baseline schedule, the unilateral decommitment schedule, and the negotiated decommitment schedule was measured.

5.2 Experimental Conditions

A standard set of experiments was designed to test the system under three operational conditions. Each set of experiments was run on the system in three separate configurations: 1) the baseline system, that is, negotiation without decommitment, 2) the system with decommitment enabled, but without negotiated decommitment, and 3) the system with full negotiated decommitment. These conditions test the first two hypotheses, that decommitment would improve goal achievement, and that negotiated decommitment would have an even more beneficial effect. Within each set of experiments, the number of targets, the number of agents, and the speed of the target were varied, to increase the number of constraints on the tracking task, and thus test the third hypothesis, that of graceful degradation of system performance under increasing constraints. The number of agents, or sensors, used was either four, six, eight or twelve. The number of targets was one, two, or three, and the target speed was set at slow (0.1 units per second), medium (0.5 units per second), or fast (1.0 units per second). Under each experimental condition, each of the performance

evaluation criteria were recorded for comparison across conditions. Additionally, tracking accuracy was measured, though not used as the main measurement of goal achievement.

For the sake of completeness, all permutations of target speed, number of targets, and number of agents were run for each of the three software configurations. How constrained each condition was, was estimated as the product of the ratio of targets to sensors and the speed of the target. In the least constrained condition, it is likely that twelve sensors tracking one target at low speed is overkill. At the other end of the spectrum, four sensors tracking three targets at high speed is extremely unlikely to be successful. Running all the conditions, however, provides a good picture of the system performance under increasing constraints, and tests the hypothesis of graceful degradation.

Agents	Targets	Target Speed	Constraint Value	Agents	Targets	Target Speed	Constraint Value
12	1	0.1	0.0083	8	1	1	0.1250
8	1	0.1	0.0125	8	2	0.5	
6	1	0.1	0.0167	12	3	0.5	
12	2	0.1		6	1	1	0.1667
4	1	0.1	0.0250	6	2	0.5	
8	2	0.1		12	2	1	
12	3	0.1	0.0333	8	3	0.5	0.1875
6	2	0.1	0.0375	4	1	1	0.2500
8	3	0.1	0.0417	4	2	0.5	
12	1	0.5	0.0500	6	3	0.5	
4	2	0.1		8	2	1	
6	3	0.1	0.0625	12	3	1	0.3333
8	1	0.5	6	2	1		
4	3	0.1	0.0750	4	3	0.5	0.3750
6	1	0.5	0.0833	8	3	1	
12	1	1		4	2	1	0.5000
12	2	0.5		6	3	1	
4	1	0.5	0.1250	4	3	1	0.7500

Table 5.1: Experimental Variables and Level of Constraint

In order to eliminate potential bias from sensor position and orientation, and target location and path, across the different conditions, these variables were held constant within each configuration. When varying the number of agents, and thus the number of sensors, sensor position and orientation was defined as shown in figures 5.1 through 5.4.




Figure 5.1: Sensor Position and Orientation with Four Agents




Figure 5.2: Sensor Position and Orientation with Six Agents




Figure 5.3: Sensor Position and Orientation with Eight Agent




Figure 5.4: Sensor Position and Orientation with Twelve Agents

The circles in the sensor position diagrams represent the location of each of the sensors within the experimental area, or “room”, which has a dimension of 80 x 60 units. The lines attached to each circle indicate its sectors, and the direction or orientation of those sectors. Across all conditions, independent of the number of agents, the defining distance for neighbors was set at 30 units.

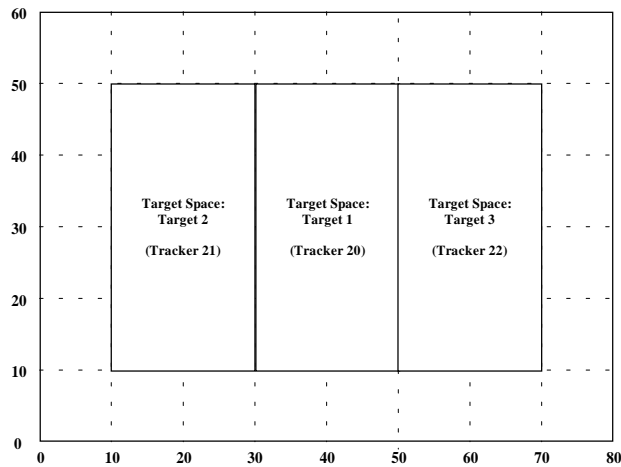


Figure 5.5: Target Position

Also independent of the number of agents, the area defined for the target path(s) was also held constant across varying numbers of targets. Figure 5.5 shows the space in the room allocated for each target. As indicated, under the experimental condition with one target, its path was always in the center space. With two targets, the paths were in the far left and right positions. Finally, with three targets, the paths occupied all three target spaces. Defining the target path spaces in such a manner allows for a somewhat equalized distribution of sensors which are capable of detecting each

target, creating less bias across conditions when assessing balanced measurements across multiple targets.

Finally, the path of each target in multiple target conditions was held constant to eliminate bias from different paths. Each target had a predefined path (scaled to fit within the space allowed in the experimental “room”), as shown in figure 5.6. The target path chosen for these experiments was an oval shape, since there are no straight line projections, and it would provide a more challenging condition for the tracker and the radar model.




Figure 5.6: Curved Target Path: Oval

Chapter 6 – Results and Analysis

The first two hypotheses of this research were concerned with the improvement of overall goal achievement by first allowing decommitment, represented by the unilateral decommitment experimental condition, and then by adding negotiated decommitment. The first four performance evaluation criteria, along with the average tracking error were used to evaluate that goal achievement.

The third hypothesis was that of graceful degradation of performance with increasing complexity of the experiment parameters, as described previously, at increasing levels of constraint. That is, the performance evaluation criteria may show decreasing performance as the level of constraint of the condition increases, but this decrease will be linear rather than exponential. The following sections describe the analysis of these measures.

5.1 Overall Goal Achievement

5.1.1 Number of Planned Measurements per Target

Figure 6.1 shows the average number of planned measurements per target across all experiments within each experimental condition. As was predicted by the first two

hypotheses, more measurements were planned in the unilateral decommitment condition than in the baseline condition. Likewise, more measurements were planned in the negotiated decommitment condition than in the unilateral decommitment condition. The difference between the baseline condition and the unilateral decommitment condition does not show statistical significance ($p=0.19$), however the negotiated decommitment shows significant improvement over both the baseline condition ($p<0.001$) and the unilateral decommitment condition ($p<0.01$).




Figure 6.1: Average Number of Planned Measurements Across all Experiments within an Experimental Condition

5.1.2 Number of Times Three or More Measurements were Taken in a Two Second Window per Target

Figure 6.2 shows the average number of times more than three measurements were taken a two second window for a given target across all experiments within an

experimental condition, allowing the `tracker` software a better opportunity to estimate target location and velocity. As with the planned measurements, the unilateral decommitment experimental condition shows an increase in grouped measurements over the baseline condition, and the negotiated decommitment condition shows an increase over the unilateral decommitment condition. Both the unilateral decommitment and negotiated decommitment conditions showed a significant increase over the baseline condition ($p < 0.001$, $p < 0.001$). However, the increase in measurements for the negotiated decommitment condition over the unilateral decommitment condition did not demonstrate significance ($p = 0.16$).

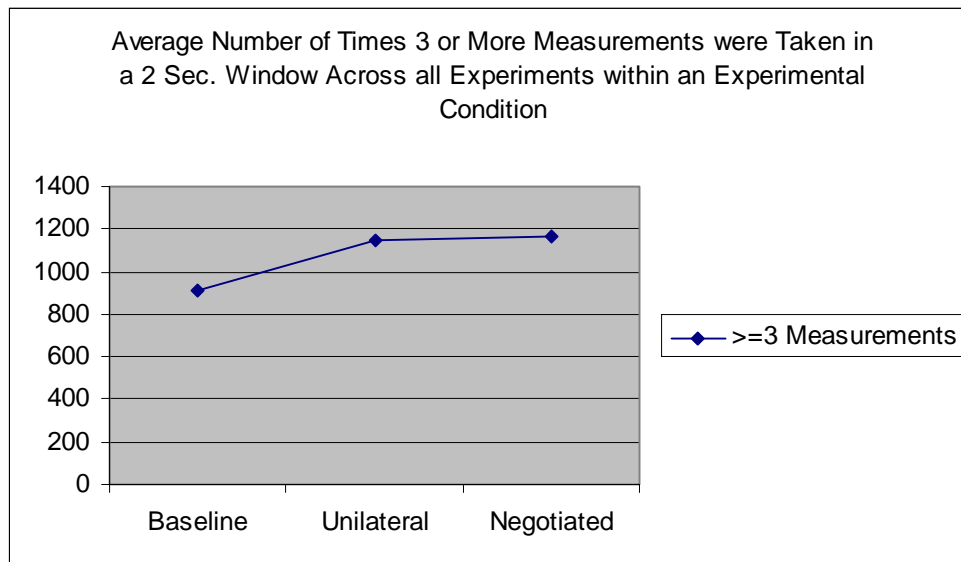


Figure 6.2: Average Number of Times 3 or More Measurements were Taken in a Two Second Window Across all Experiments within an Experimental Condition

5.1.3 Balanced Measurements Across Multiple Targets

To measure the balance of measurements across targets in conditions where multiple targets were present, the standard deviation between the number of measurements of different targets in the same experimental run was used to represent the amount of difference. Thus, the smaller the number, the more balanced the measurements across targets. Figure 6.3 depicts this balance of measurements by experimental condition. Unexpectedly, the baseline condition shows a better balance of measurements than does either the unilateral decommitment condition or the negotiated decommitment condition, however neither of these differences showed significance ($p=0.12$, $p=0.20$). The negotiated decommitment condition does show a slight improvement over the unilateral decommitment condition, but again the difference was not significant ($p=.45$).




Figure 6.3: Balance of Measurements across Multiple Targets




Figure 6.4: Average Total Number of Measurements Across all Experiments within an Experimental Condition

5.1.4 Total Number of Measurements Taken

Figure 6.4 shows the average total number of measurements taken per target by experimental condition. As with the first two performance evaluation criteria, unilateral decommitment showed significantly more total measurements than the baseline condition ($p < 0.001$), the negotiated decommitment condition showed more measurements than the baseline condition ($p < 0.001$), and the negotiated decommitment condition showed only a slight increase over the unilateral decommitment condition ($p = 0.16$).

5.1.5 Average Tracking Error

Figure 6.5 shows the average tracking error across experimental conditions.

Unilateral decommitment shows slight improvement over the baseline condition, however the improvement was not significant ($p=0.49$), the negotiated decommitment condition shows slight improvement over the baseline condition ($p=0.30$), and the negotiated decommitment condition shows very slight improvement over the unilateral decommitment condition ($p=0.31$).




Figure 6.5: Average Tracking Error Across all Experiments within an Experimental Condition

5.1.6 Discussion

All performance evaluation criteria showed goal achievement improvement trends as predicted by the hypotheses, with the exception of the balance of measurements across multiple targets. The magnitude of performance improvement between the negotiated decommitment and unilateral decommitment experimental conditions as shown by the evaluation criteria was not as large as expected, however. In order to explore this in more detail, the utility of making or breaking a commitment was measured in each experimental run for each of the experimental conditions to determine what the decision would have been under that condition. The decision made by the agent, however, was the one dictated by the experimental condition. Table 6.6 shows the agents' perceived average utility per millisecond across all runs for each of the experimental conditions.

	Average Utility per Millisecond
Existing Schedule	5.518
Baseline Schedule	6.228
Unilateral Decommitment Schedule	6.297
Negotiated Decommitment Schedule	6.299

Table 6.1: Average Utility per Millisecond

The consistent increase in perceived utility in a schedule through the experimental conditions indicates that agents do, in fact, make rational decisions, and these

decisions are perceived to have higher utility in the unilateral and negotiated decommitment conditions. However, the greatest increase in perceived utility occurs between the existing schedule and the baseline condition, or the agents capability to make commitments versus not making commitments. Allowing an agent to drop commitments results in a somewhat higher perceived utility, and negotiating the decommitment only adds a slight increase.

An explanation for this minimal increase in perceived utility can be found in the way the `tracker` and `radar model` programs provide predictions for future target locations. The `tracker` program provides an estimate of a target's location and velocity at a given point in time. This estimate is then passed to the `radar model` which projects the target moving in a straight line trajectory from that location, at that velocity. The `radar model` then calculates the entry and exit of the target from each sensor's visibility area. An agent initiates negotiations with its neighbors based on this projection.

Even if the `tracker` provided perfect estimates, a target that changes direction will no longer be traveling in the initially projected direction. Thus any commitments made by neighbors based on that initial projection may no longer have value. This projection error is increased when the `tracker` estimate is not correct. In the unilateral condition, agents drop commitments when they have lower utility and a task of higher utility presents itself, and thus achieve an increase in perceived utility.

In the negotiated decommitment condition, an agent negotiates with the original requesting agent before dropping the commitment. In the vast majority of negotiated decommitments, the originating agent agrees with the decommitting agent that the task is no longer of value, and the commitment is dropped. In this case, the perceived utility becomes equal to that of a unilateral decommitment. The few cases where the originating agent does not agree to allow the decommitting agent to drop the commitment accounts for the very slight increase in perceived utility in the negotiated decommitment condition over the unilateral decommitment condition.

As an illustration of the difference between the potential for utility increase between the unilateral and negotiated decommitment conditions, an example of a successful counter offer to a decommitment occurred in an experimental run with 6 agents, 1 target, at medium speed. In this particular case, a neighbor requested that a task be performed for 36 seconds. The agent assessed the requested task's utility at 6.387. Under the baseline condition, the utility for the duration of the schedule that the task would have spanned was 6.309, a lower utility than the requested task. But because one or more tasks in that span were commitments to other neighbors, the agent would have been unable to schedule that task. Under the unilateral decommitment condition, the agent would have been able to schedule that task, and the resulting utility for that time span would have been 6.387, the utility of the requested task. Under the negotiated decommitment condition, however, the agent negotiated with the originating agent, and the originating agent re-evaluated the current conditions and

returned a counter offer. The agent considering decommitment then reassessed the utility of the commitment it had considered dropping, and the newly assessed utility was 8.907. Thus, under the negotiated decommitment condition, the requesting agent's request was refused, and the utility assessment of the existing schedule increased because of the new assessment of the existing commitment.

Unfortunately, because of erroneous target location projections, the conditions for higher counter offers to a negotiated decommitment don't occur often enough to demonstrate a larger benefit, either in perceived utility or in the performance evaluation criteria, for negotiated decommitment over unilateral decommitment. The slight improvement in perceived utility between experimental conditions is reflected in the actual experimental results. That is, the increase in overall goal achievement between the unilateral and negotiated decommitment conditions does not show as great a difference across conditions as was expected.

5.2 Graceful Degradation of Performance

The third hypothesis predicted that as experimental conditions were increasingly constrained, graceful degradation of performance would be shown, rather than catastrophic degradation. The level of constraint on conditions was defined as the product of the target speed and the ratio of sensors to targets. Figure 6.6 shows the average tracking error across the defined levels of constraint.




Figure 6.6: Average Tracking Error Across Levels of Constraint

With increasing levels of constraint, although the average tracking error does not show a smooth pattern, it does indicate that tends to stabilize at higher levels, rather than failing catastrophically. However, the irregular pattern implies that the definition of constraint levels was not accurate. With further investigation, an interesting phenomenon was uncovered. Contrary to expectation, tracking error actually decreases as the target speed increases, as shown in Figure 6.7. This phenomenon occurs across all of the experimental conditions approximately equally.

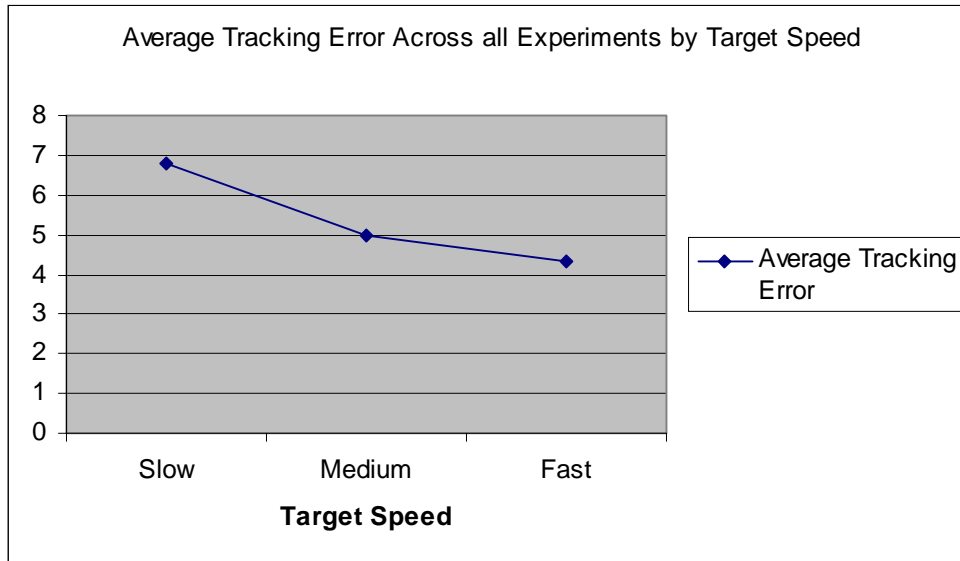


Figure 6.7: Average Tracking Error Across All Experimental Conditions by Target Speed

Since the `tracker` program relies on changes in amplitude and frequency measurements submitted to it by sensors, it appears that at faster target speeds, the additional changes improve the tracking performance, and at slower speeds the measurements don't change often enough to provide good tracking results.

Removing the target speed factor from levels of constraint in the experiments, and defining levels of constraint as only the ratio of targets to sensors results in the tracking error performance shown in Figure 6.8.




Figure 6.8: Average Tracking Error Across all Experiments by Target to Sensor Ratio

This shows a more linear increase in tracking error as the ratio becomes higher, thus supporting the hypothesis of graceful degradation with increasing constraints across all experimental conditions. Configurations with a higher target to sensor ratio than 3:4 were not tested in this research, therefore it can be stated only that the overall system exhibited graceful degradation in tracking error up to this limit.

Individual performance evaluation criteria also show a linear or near linear degradation in performance as the target to sensor ratio increases, however there is little difference between experimental conditions. Figure 6.9 shows the number of planned measurements across increasing constraints, by each experimental condition.

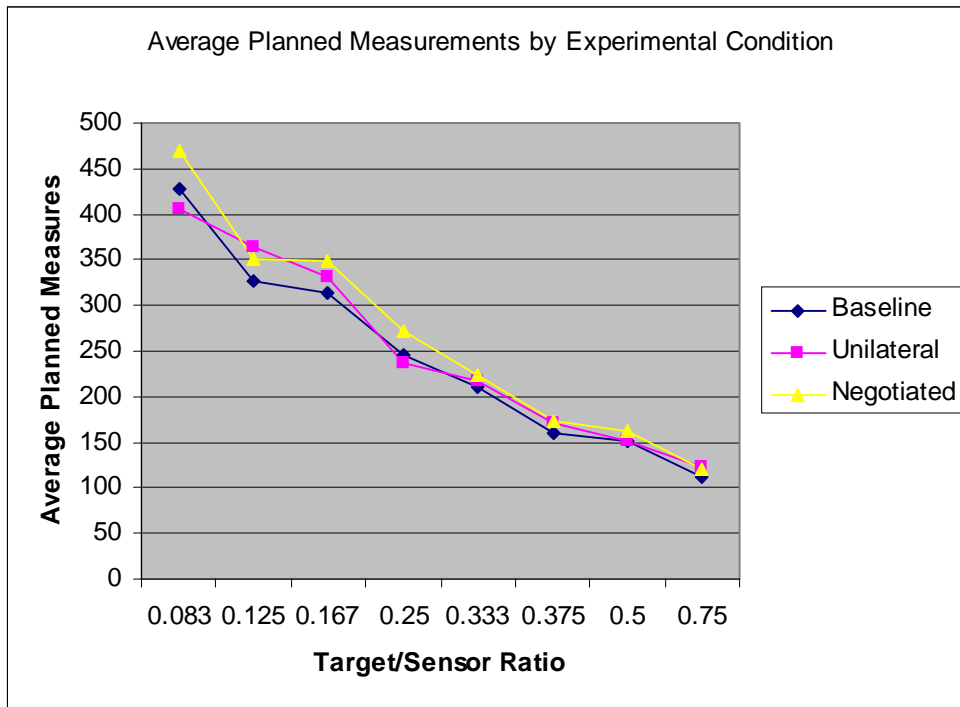


Figure 6.9: Average Planned Measurements by Experimental Condition

The number of planned measurements decreases with increasing constraints in a fairly linear manner, regardless of experimental condition.

Figure 6.10 shows the number of times three or more measurements were taken within a two second window by constraint, and by experimental condition. The unilateral and negotiated decommitment experimental conditions appear to level off at a somewhat higher level at more highly constrained conditions than does the baseline condition, but without more highly constrained conditions it is impossible to tell if this trend would continue.




Figure 6.10: Average Times Three or More Measurements were Taken in a Two Second Window

Figure 6.11 depicts the balance of measurements across multiple targets by level of constraint for each of the experimental conditions. Of all the performance evaluation criteria, this shows the least graceful degradation in performance. With more targets per sensor, the balance becomes increasingly worse, with negotiated decommitment showing the highest difference between measurements of targets.




Figure 6.11: Averaged Standard Deviation of Measures Across Multiple Targets

The average total number of measurements per target shows a similar pattern as that of the number of times three or more measurements were taken in a two second window. That is, both unilateral and negotiated decommitment appear to level off at a slightly higher level than does the baseline condition, at higher levels of constraint. This is shown in Figure 6.12.




Figure 6.12: Average Total Number of Measurements per Target

Finally, the average tracking error by experimental condition shows a linear trend in degradation over increasing constraints for each of the conditions, very similar to that of the overall system degradation in performance discussed at the beginning of this section. The average tracking error by experimental condition is shown in Figure 6.13.




Figure 6.13: Average Tracking Error Across Increasing Constraints

It was hypothesized that the decommitment conditions would exhibit graceful degradation with increasing constraints, and with the exception of the balance of measurements across multiple targets, this has been demonstrated. However, the baseline condition, without decommitment, also shows graceful degradation across increasing constraints, and there appears to be little discernable difference between the experimental conditions.

Chapter 7 – Conclusions and Future Directions

This research investigated the decommitment of agents in a collaborative and benevolent society of agents. Andersson and Sandholm (1998) call these “social welfare” or SW agents, in which each agent attempts to maximize the profit of the whole society, as opposed to self-interested agents that attempt to maximize their own individual profit. Review of existing literature shows that decommitment has been investigated primarily in the context of self-interested agents, and not in cooperative agents societies. Furthermore, the concept of negotiated decommitment has not been at all in the literature. The contribution of this research, then, is to demonstrate that intentional, or rational, decommitment improves goal achievement in cooperative and benevolent agent societies, and that negotiated decommitment provides even more benefit than unilateral decommitment, despite the additional communication overhead.

In this research, agents agreed to commitments based on their assessment of the utility of taking on that commitment, or the commitment’s *value*. Decommitment was only considered when alternative tasks provided a higher perceived utility than a current commitment. Even when alternative tasks provided a higher utility than an existing commitment, the commitment was not dropped unless the commitment *strength* indicated that dropping the commitment would still provide benefit to the overall system. Unilateral decommitment considered only the decommitting agents

own perception of commitment value and strength, while negotiated decommitment requested input from agents that would be affected by the decommitment.

The hypotheses of this research are that given a society of agents that are autonomous, rational, collaborative, benevolent, and that are capable of negotiation and planning time-dependent activities:

1. Allowing agent decommitment will improve the overall goal achievement of the system.
2. When an agent's decommitment to a task or goal affects other agents, negotiating the terms of the decommitment with those other agents will be more beneficial than unilateral decommitment, despite the associated increase in overhead. The improvement is expected as a result of a more refined estimation of global utility when more than one agent contributes to that estimate.
3. As the number of constraints on the system increase, the overall goal achievement of the system will degrade gracefully.

Results indicated that both unilateral and negotiated decommitment performed as hypothesized, that is, overall performance improved. Negotiated decommitment provided additional benefit over the performance of unilateral commitment. Graceful degradation of performance was demonstrated in all conditions as the constraints on the system increased, although there was little difference across the experimental conditions. Under all conditions, agents demonstrated rationality in decision making

and both negotiated decommitment and unilateral decommitment provided a higher locally assessed utility than did the baseline condition.

Although support for the research hypotheses was shown in the experimental results, the magnitude of increase between the unilateral and negotiated decommitment conditions was not as large as expected. The domain on which the research was tested may not have been ideal for demonstrating the potential of negotiated decommitment for several reasons. First, as was discussed previously, making commitments based on erroneous future predictions can reduce the benefit of negotiated decommitment, since it is almost always more beneficial to drop a commitment than to keep it. A different domain which has more reliable predictions of future conditions may prove to be a better testing ground than the multi-sensor target tracking domain.

Second, another difficulty with the particular target tracking domain used was the communication bottleneck. All inter-agent communications were required to be passed through the radar simulator, `Radsim`, creating a severe bottleneck and resulting in a large amount of timed out negotiations and requests that were received too late to be considered. If the current domain were to be modified, direct agent to agent communication would be highly desirable

There are several avenues open for additional research in this area. The most obvious is that of investigating the effects of negotiated decommitment on domains with

different characteristics. Another interesting area of investigation is that of the commitment *value* and *strength* measures. Sensitivity testing of the factors in these equations may yield a more accurate estimate of local utility, and thus better decision making by the agent. In the area of multi-sensor target tracking, the phenomena of improved tracking performance at increased target speeds raises the question of what the limits may be. One would expect that at some higher speed, tracking performance would degrade, and the parameters that determine what this limit is provide an interesting research question. Finally, the question of graceful degradation of performance could be investigated under conditions that show higher levels of constraint than what was addressed in this research.

Overall this research has demonstrated that both unilateral and negotiated decommitment performed as hypothesized, that is, overall performance improved, and that negotiated decommitment provided additional benefit over the performance of unilateral commitment. Graceful degradation of performance was also demonstrated in all conditions as the constraints on the system. Finally, under all conditions, agents demonstrated rationality in decision making and both negotiated decommitment and unilateral decommitment provided a higher locally assessed utility than did the baseline condition.

References

- Andersson, M.R. and Sandholm, T.W. 1998. "Leveled commitment contracts with myopic and strategic agents", *AAAI/IAAI*, pp. 38-45.
- Booch, G. 1994. *Object-Oriented Analysis and Design with Applications*, Addison-Wesley, Reading, MA.
- Caplow, T. 1956. "A theory of coalitions in the triad", *American Sociological Review*, 21, pp. 489-493.
- Castelfranchi, C. 1998. "Modeling social action for AI agents", *Artificial Intelligence*, 103 (1-2), pp. 157-182.
- Conry, S.E., Meyer, R.A and Lesser, V.R. 1988. "Multistage negotiation in distributed planning", in Gasser and Bond (Eds.), *Readings in Distributed Artificial Intelligence*, Morgan Kaufmann, San Mateo, CA, pp. 367-384.
- Conte, R., Miceli, M. and Castelfranchi, C. 1991. "Limits and levels of cooperation: disentangling various types of prosocial interaction", in Demazeau, Y. and Muller, J.-P. (eds.), *Decentralized A.I. Vol. 2*, pp. 147-157.
- Decker, K.S. 1987. "Distributed problem-solving techniques: a survey", *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-17 (5), pp. 729-740.
- Durfee, E.H. and Lesser, V.R. 1991. "Partial global planning: a coordination framework for distributed hypothesis formation", *IEEE Transactions on Systems, Man, and Cybernetics*, 21(5), 1167-1183.
- Excelente-Toledo, C.B., Bourne, R.A. and Jennings, N.R. 2001. "Reasoning about commitments and penalties for coordination between autonomous agents", *Proc. 5th Int Conf on Autonomous Agents (Agents-2001)*, Montreal, Canada, pp. 131-138.
- Faratin, P., Sierra, C., and Jennings, N.R. 1998. "Negotiation Decision Functions for Autonomous Agents", *Int. Journal of Robotics and Autonomous Systems*, 24 (3-4), pp. 159-182.
- Fitoussi, D. and Tennenholtz, M. 2000. "Choosing social laws for multi-agent systems: Minimality and simplicity", *Artificial Intelligence*, 119 (1-2), pp. 61-101.

Gamson, W. 1961. "A theory of coalition formation", *American Sociological Review*, 21, pp. 373-382.

Jennings, N.R. 2000. "On agent-based software engineering", *Artificial Intelligence*, 117 (2), pp. 277-296.

Kraus, S., Sycara, K. and Evenchik, A. 1998. "Reaching agreements through argumentation: a logical model and implementation", *Artificial Intelligence*, 104 (1-2), pp. 1-69.

Larson, K.S. and Sandholm, T.W. 1999. "Anytime coalition structure generation: an average case study", *Autonomous Agents '99*, Seattle, WA, pp. 40-47.

Lander, S.E. and Lesser, V.R. 1993. "Understanding the role of negotiation in distributed search among heterogeneous agents", *Proceedings of the 12th International Workshop on DAI*, Hidden Valley, PA.

Meyer, J.-J.Ch., van der Hoek, W. and van Linder, B. 1999. "A logical approach to the dynamics of commitments", *Artificial Intelligence*, 113 (1-2), pp. 1-40.

Mouaddib, A.-I. 1997. "Progressive negotiation for time-constrained autonomous agents", *Autonomous Agents '97*, Marina Del Rey, CA, pp. 8-15.

Newell, A. 1982. "The knowledge level", *Artificial Intelligence*, 18, pp. 87-127.

Noriega, P., & Sierra, C. 1996. "Towards layered dialogical agents", *Proceedings of the 3rd ECAI Workshop on Agents Theories, Architectures and Languages (ATAL'96)*, Budapest, Hungary, 157-171.

Norman, T.J., Sierra, C. and Jennings, N.R. 1998. "Rights and commitments in multi-agent agreements", *Proc. 3rd Int. Conf. on Multi-Agent Systems (ICMAS-98)*, Paris, France, pp. 222-229.

Panzarasa, P. and Jennings, N.R. 2001. "Negotiation and joint commitments in multi-agent systems", *Proc. 2nd Int Workshop on Modeling Artificial Societies and Hybrid Organisations*, Vienna, Austria, pp. 1-63.

Parsons, S. & Jennings, N. R. 1996. "Negotiation through argumentation—a preliminary report", *Proceedings of the International Conference on Multi-Agent Systems*, Kyoto, Japan.

Parsons, T.S., Sierra, C. and Jennings, N.R. 1998. "Agents that reason and negotiate by arguing", *Journal of Logic and Computation*, 8 (3), pp. 261-292.

Raiffa, H. 1982. *The Art and Science of Negotiation*, Cambridge, Mass., Harvard University Press.

Rao, A.S. 1995. "Integrated agent architectures: execution and recognition of mental states", in Chengqi Zhang, Dickson Lukose (Eds.): *Distributed Artificial Intelligence: Architecture and Modelling, First Australian Workshop on DAI*, Canberra, ACT, Australia, pp. 159-173.

Rao, A.S. and Georgeff, M.P. 1995. "BDI Agents: from theory to practice", Proceedings of the 1st International Conference on Multi-Agent Systems (ICMAS '95), June 12-14, San Francisco, CA, pp. 312-319.

Robinson, W.N. and Volkov, V. 1998. "Supporting the negotiation life cycle", *Communications of the ACM*, 41 (5), pp. 95-102.

Rosenschein, J.S. and Zlotkin, G. 1994. "Designing conventions for automated negotiations", *AI Magazine*, 15 (3), pp. 29-46.

Sandholm, T., Larson, K., Andersson, M., Shehory, O. and Tohmé, F. 1999. "Coalition structure generation with worst case guarantees", *Artificial Intelligence*, 111 (1-2), pp. 209-238.

Sandholm, T.W. and Lesser, V.R. 1995. "Coalition formation among bounded rational agents", *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI-95)*, Montreal, Canada, pp. 662-669.

Sandholm, T.W. and Lesser, V.R. 1996. "Advantages of a leveled commitment contracting protocol", *Proceedings of the National Conference on Artificial Intelligence*, pp. 126-133.

Schut, M. and Woolridge, M. 2001. "Principles on intention reconsideration", *AGENTS '01*, May 28 – June 1, Montreal, Quebec, Canada, pp. 340-347.

Sen, S. and Durfee, E. 1994. "The role of commitment in cooperative negotiation", *International Journal on Intelligent Cooperative Information Systems*, 3(1), pp. 67-81.

Sen, S. and Durfee, E. 1996. "A contracting model for flexible distributed scheduling", *Annals of Operations Research*, 65, pp. 195-222.

- Sen, S. and Durfee, E. 1998. "A formal study of distributed meeting scheduling", *Group Discussion and Negotiation*, 7, pp. 265-289.
- Shehory, O. and Kraus, S. 1998. "Methods for task allocation via agent coalition formation", *Artificial Intelligence*, 101 (1-2), pp. 165-200.
- Shenoy, P.P. 1978. "On coalition formation in simple games: a mathematical analysis of Caplow's and Gamsen's theories", *Journal of Mathematical Psychology*, 18 (2), pp.177-194.
- Shenoy, P.P. 1979. "On coalition formation: a game theoretic approach", *Journal of Game Theory*, 8 (3), pp.133-164.
- Shoham, Y. and Tennenholtz, M. 1992. "On the synthesis of useful social laws for artificial agent societies", *Proceedings of AAAI*, pp. 276-281.
- Smith, R.G. 1980. "The contract net protocol: high-level communication and control in a distributed problem solver", *IEEE Transactions on Computers*, 29 (12), pp. 1104-1113.
- Soh, L.-K. and Tsatsoulis, C. 2001a. "Logical negotiation protocol and its implementation, *Artificial Intelligence*, under review.
- Soh, L.-K. and Tsatsoulis, C. 2001b. "A dynamic negotiation-based coalition formation in a cooperative multiagent system with incomplete information", unpublished draft.
- Tohme, F. and Sandholm, T. 1999. "Coalition formation processes with belief revision among bounded-rational self-interested agents", *Journal of Logic and Computation*, 9 (6), pp. 793-815.
- Woolridge, M. and Jennings, N.R. 1995. "Intelligent agents: theory and practice", *The Knowledge Engineering Review*, 10 (2), 114-152.
- Xing, J. and Singh, M.P. 2001. "Formalization of commitment-based agent interaction", *SAC 2001*, Las Vegas, NV., pp. 115-120.
- Zlotkin, G. and Rosenschein, J.S. 1991. "Cooperation and conflict resolution via negotiation among autonomous agents in noncooperative domains", *IEEE Transactions on Systems, Man, and Cybernetics*, 21 (6), pp. 1317-1324.