

## Chapter 4: Multiagent negotiation models for power system applications

J. McCalley, Z. Zhang, V. Vishwanathan, and V. Honavar

### 4.1 Introduction

There are a wide range of power system decision problems, traditionally falling under one of the three categories of operations, maintenance, and planning, with the delineation between categories derived from the nature of the decision and the time horizon. Some of these decision problems include generation dispatching, fuel scheduling, control-room preventive and corrective actions, incident restoration, transmission service scheduling, unit commitment, transmission equipment maintenance, control system planning, and transmission upgrades. All of these share common attributes, among which are:

- Resulting decisions have system-wide impact and therefore require significant coordination of information among the various decision-makers.
- Higher system integrity is only achieved with greater allocation of financial resources.
- The essential decision variables are inherently uncertain.
- Uncertainty is reduced via acquisition and processing of information
- Important information tends to be spatially dispersed.
- Complex and computationally-intensive applications are required.

As a result, decision-making support aids require modeling of multiple objectives, application of significant computational resources, and use of flexible data access capability. Mathematical programming has been and continues to be a mainstay for such decision problems. However, traditional optimization tools, by their very nature, assume the existence of a single and benevolent decision-maker that has centralized access to all information, and coordination between decision-makers is embedded in the processes and of no threat to individuals carrying out those processes. Such was the case in the traditionally regulated world where ownership of all facilities, access to all information, and all authority for decision rested within the single umbrella of the vertically integrated utility company. With the advent of industry restructuring and associated organizational disaggregation, however, facility ownership is heavily balkanized, and in-

formation access and decision-making authority is quite limited for any one particular organization. Even more, the various organizations comprising the industry are not necessarily cooperative one with another; in fact, many portions of the restructured industry are intentionally organized to be competitive. Yet the need for coordinated decision remains, because the operational integrity of the system and individual components demands it. This necessitates a new paradigm to build upon and ultimately replace the centralized decision approach, enabling optimized decisions in an environment of highly distributed information and a multiplicity of competing stakeholders.

The industry has attempted to retain decision-making ability using traditional optimization tools, but it has come at the expense of forming new, centralized and competitively neutral authorities such as independent system operators (ISOs) and reliability authorities (RAs) to coordinate system operations and issues related to system reliability. These organizations arbitrate those decisions where conflict between two or more parties may otherwise arise. For example, during operationally stressed conditions having excessive risk of load interruption, a centralized authority generally selects the units to redispatch in order to mitigate the risk. Another example is maintenance: given requests for simultaneous maintenance outage of multiple components (generators, lines, and/or transformers) such that network integrity is excessively compromised, a centralized authority generally determines the sequence and timing of the maintenance tasks. In both of these cases, a conflict (which generators to redispatch in case 1 and which maintenance tasks to postpone in case 2) is settled by the arbitration of the central authority. The technology which motivates this book, multiagent systems (MAS), may offer a viable alternative to this arrangement, or at least a useful complement, through the use of software agents equipped with negotiated decision-making capabilities operating within a MAS so as to coordinate decision-making of competing stakeholders. In multiagent negotiation systems (MANS), stakeholders, represented by agents, engage in negotiation, proposing and counter-proposing until an outcome is identified that is satisfactory to all. Here, a software agent<sup>1</sup>, armed with a coded negotiation model, represents each stakeholder, and

---

<sup>1</sup> It is important to clarify terms at the beginning. A *stakeholder* represents a human individual or human organization that has interest in the stated decision-making problem would be one of the decision-makers if allowed to participate. Later, we use the term *player* and *party* to more explicitly refer to a stakeholder involved in a human-to-human negotiation. On the other hand, an *agent* is first and foremost a software entity, secondly, one that satisfies the usual criteria for agency (e.g., a computer system situated in some environment capable of autonomous action to meet its design objectives [1]), and thirdly, for our purposes here, one that has the essential social ability of communicating (sending and receiving messages). Thus, a stakeholder may be represented by a party, or by an agent. A party always represents a stakeholder. Although an agent does not necessarily have to represent a stakeholder (there may exist "functional" agents, for example), we assume in this chapter that an agent does. In effect then, within the domain of this chapter, "agent" is the software encapsulation of "party."

conflict resolution is achieved via inter-agent message exchange until agreement is reached. MAS is an essential enabling technology because it provides the necessary infrastructure in terms of model instantiation and maintenance together with the communication needs, including messaging, directory services, and communication protocols [1].

There are at least two distinct ways in which the power industry will benefit from successful implementation of MANS: better decisions and better models. Better decisions may be expected because: (a) The ability to perform computer-evaluation of relevant issues, including accessing complex, distributed data, is inherent to the negotiation itself. This is in contrast to human-based negotiation where obtaining additional information or performing additional processing is typically done outside the time and space given to the negotiation process. (b) The negotiation speed is significantly increased. In contrast to human-based negotiation where negotiation speed depends on the limitations of the human negotiators, computer-based negotiated decisions may be reached as fast as network bandwidth and computer processing power allow. This not only provides for enhancing existing negotiated decision-making scenarios but also introducing negotiated decision-making where it was previously thought to be untenable. For example, networked negotiation enables consideration of negotiated decision-making between control centers, even between individual generation and/or transmission companies, following outages when typically decision-time is quite short.

The second way in which the power industry will benefit is that MANS offers a modeling framework that enables study of important power industry characteristics for which good models are not presently available. One of these characteristics is the level of centralization in decision-making, i.e., negotiated decision-making allows simulation and study of decision quality as the decision framework moves from being highly centralized (i.e., vertically integrated utilities) to entirely distributed. Another characteristic that can be studied with MANS is the flow of information; in particular, one can document the movement of information as well as the amount. Such documentation offers to illuminate organization effectiveness as a function of information flow, and to study the effect of providing new information flow paths.

The rest of this chapter is organized as follows. Section 4.2 reviews literature in terms of (a) negotiation theory and (b) computer-based negotiations. Section 4.3 describes implementation issues together with a description of the infrastructure we have used in our work. Section 4.4 describes power-industry application areas. Section 4.5 concludes.

## **4.2 Negotiation theory and agents: a review**

Endowing agents with advanced social abilities, such as negotiation, for use within multiagent systems, has been of interest since the 1980's, but study of negotiation as a fundamental form of human interac-

tion has been ongoing throughout the 20<sup>th</sup> century, and an awareness of these developments is essential for understanding the recent and intimately related work in MANS. Section 4.2.1 focuses on literature from decision science, economics, and anthropology; more recent literature related to computer-based negotiation mechanisms are discussed in Section 4.2.2.

#### 4.2.1 Basics of negotiation theory

Negotiation is a fundamental form of human interaction, and we see it in labour-management disputes, international diplomacy, governmental processes, business relations, and interpersonal relations. Despite its prevalence throughout all human history, it was not until the middle of the 20<sup>th</sup> century before development of a theory for negotiation was initiated. This effort had roots in a number of different disciplines, including decision science, economic bargaining theory, social psychology, political science, industrial sociology, and social anthropology [2]. We do not attempt a comprehensive literature review here but rather provide basic concepts on which we draw in instantiating negotiation models within agents.

There are two sub-disciplines within decision science that need particular attention in order to do justice to the field of negotiation theory. The first is multi-criteria decision-making (MCDM), because it is MCDM that provides a number of different decision approaches for multi-criteria decision problems. Some of these approaches include [3] weighting methods such as analytical hierarchy process, Electre IV, goal programming, evidential theory, and utility-based approaches, where we typically search for efficient solutions  $y^*$ , i.e., those solutions for which there exist no other solutions that can outperform  $y^*$  in all criteria. Of the various MCDM approaches, it is the utility-based approaches that have had particular influence on evolution of negotiation theory. A well-known and simple decision criterion is to choose the action which maximizes the expected value of benefit. Thus, if we can associate with each course of action A a set of outcomes characterized by their benefits  $c_1, c_2, \dots, c_n$  and corresponding probabilities  $p_1, p_2, \dots, p_n$ , we desire to select the course of action that has the largest value of  $\sum p_i c_i$ . Bernoulli [4], and later von Neumann and Morgenstern [5] and others [6,7,8,9] argued that rather than using expected value, the rational way for people to evaluate decision problems is on the basis of expected utility  $EU(A) = \sum p_i u(c_i)$  where  $u(\bullet)$  is a utility function that characterizes the decision-makers preferences with respect to the possible benefits of each outcome. Construction of utility functions...

The second sub-discipline that needs particular attention is the theory of *competitive problems* [10], characterized by decision scenarios

where certain of the decision variables are controlled by two or more independent parties having different interests. This discipline, which has largely grown out of utility-based decision approaches, has formed the basis for much of the non-agent and agent-related work in negotiation theory. The most influential aspects of this work fall under the *theory of games* in which two or more *players* (i.e. parties) choose courses of action and in which the outcome is affected by the combination of choices taken collectively [2,5,6,7,8,9,10]. A key assumption is that players behave *rationally*, where rational behavior is characterized by action selection, by each party, so as to maximize individual expected utility. Additional assumptions include (a) there is a fixed set of rules that specify what courses of action can be chosen; (b) there are well-defined end-states that terminate the game; (c) associated with each end-state are player-specific payoffs; (d) all players have perfect knowledge with regard to the rules, the range of outcomes, probabilities, and payoffs, and each player's preferences; (e) there is no interference or influence from the outside world. The central question addressed is: for a specified game, under assumptions a-e, what will be the utility vector on which the players will agree? The most well-known example of such a game is the so-called prisoners' dilemma whereby the district attorney has two robbers in different cells. If both confess, both get 8 years jail time; if neither confesses, both get 5 years, and if one confesses and the other does not, the confessor gets 2 years and the other 10. Game theory provides several different models for studying player decisions.

A simple game-theoretic model is identified by Raiffa in [11], where it is assumed that by analyzing the consequences of no agreement, each party can establish a threshold value to be used for decision. Define  $x^*$  as the final-contract value, the sellers reservation price  $s$  that represents the very minimum price for which the seller will sell, the buyers reservation price  $b$  that represents the very maximum price for which the buyer will buy. Then the zone of agreement is the interval  $(s,b)$ , assuming  $s < b$ . If  $b < s$ , then agreement is not possible. The buyer's surplus is  $b - x^*$ , the sellers surplus is  $x^* - s$ , and both buyer and seller try to maximize their surplus.

Other game theoretic models pertaining to negotiation are identified in chapter 2 of [12], the most important of which is Zeuthen's two-party model. Here, if we define  $x_0$  as the demand offered by party 2, party 1 identifies its expected payoff utility from demanding  $x$  according to  $\Delta u_1 = u_1(x)[1 - p(x)] - u_1(x_0)$ , where  $p(x)$  is the probability that party 1's insistence on  $x$  will result in a walk-out by party 2 (and  $1 - p(x)$  is the probability that party 1 will agree). If  $p(x)$  is too large, then  $\Delta u_1$  could be negative implying expected utility decreases for party 1, an undesirable situation. Therefore, in order to decide whether to counter-propose with  $x$ , party 1

determines the highest  $p(x_m)$  that does not give  $\Delta u_1 < 0$ . This is found from solving  $\Delta u_1 = 0 = u_1(x_m)[1 - p(x_m)] - u_1(x_0) \Leftrightarrow p(x_m) = [u_1(x_m) - u_1(x_0)] / u_1(x_m)$ . Thus, party 1 should counter-propose a certain value  $x$  only if  $p(x) < p(x_m)$ . Party 2 is assumed to use a similar model.

Cross also provides a more advanced two-party model where time and money are the issues. Each party knows their own desired settlement,  $q_1$  and  $q_2$ , respectively. Also, each party has an expectation of their opponent's rate of concession (change in offer per time period); party 1 uses  $r_2$  and party 2 uses  $r_1$ . Then there is a cost to each party for each time interval that passes without agreement,  $C_1$  for party 1 and  $C_2$  for party 2. There is also a discount factor to express the money in terms of present value, so that the present value cost is given by  $Z = (C_1/a)(1 - e^{-aT})$  where  $a$  is the discount factor and  $T$  is the expected time to reach agreement given by  $T = (q_1 + q_2 - M) / r_2$ . Then the total present value to party 1 of targeting a desired settlement is given by  $U_1 = u_1(q_1) - Z$ . Cross identifies increasingly rigorous ways to update  $r_2$  (learning) based on the negotiation process.

One limitation to game theory is that it is pre-occupied with outcome, discussed in terms of equilibria (e.g., Nash, perfect, dominant), rather than the process (or mechanism) used to arrive at that outcome. This point is central to the goal of automating multi-party decision-making because we need the capability of implementing the mechanism to achieve this goal. Thus, we turn to the closely related negotiation theory.

There are at least six different mechanisms of reaching a collective decision among 2 or more parties, including persuading, educating, manipulating, coercing, appealing to an authority, and negotiation [13]. Of these, the last two, arbitration and negotiation, are two distinctive forms of identifying agreements between two or more parties that have significantly more formality and structure. Arbitration provides a mechanism which selects a single outcome as the point of agreement between the parties. Judges often assume this role in legal disputes, and, as observed in Section 4.1, so do certain kinds of power system decision-making authorities. In arbitration, the parties direct their communication towards a third party, but not to each other. Arbitration, with only a single decision-maker, is most effective when parties seek to agree over values, norms, and the assessment of facts. Negotiation, on the other hand, is the *joint-decision* process of forming and revising offers, by each involved party, whereby offers are made with the intention to converge to an agreement, without the presence of a third-party decision-maker. (Reference [2] argues that this definition applies more appropriately to *bargaining* with a broader definition used for negotiation that includes the initiation and recognition of the motivating need, the process, the final outcome, and the execution of that outcome). In negotiation, the focus of communication is (are) the other

party (parties). Generally, negotiation is performed as a result of a conflict or dispute between two or more parties, and the negotiation objective is to resolve the dispute. Negotiation is most effective in a situation of scarcity when parties seek the same resources without there being enough to satisfy both [2]. These types of negotiations have been characterized as either *strident antagonist* or *cooperative antagonist* [11]. The former is characterized by completely distrustful and malevolent (towards one another) parties, as would be the case when authorities negotiate with kidnappers or airline hijackers. The latter is characterized by entirely self-interested and disputing parties but ones that recognize and abide by whatever rules exist. A third type of negotiation is called *fully cooperative* [11], where the parties have different needs, values, and opinions, but they share information, expect total honesty, perform no strategic posturing, and think of themselves as a cohesive entity with intention to arrive at the best decision for the entity, as would be the case for a happily married couple. We are interested here in the two different levels of cooperative negotiations, since they better typify the various types of power system decisions problems. For example, a negotiation involving two transmission owners over equipment maintenance schedules is a good example of cooperative antagonists. A negotiation involving two ISOs sharing responsibility for operational integrity of different portions of the network in the same interconnection, over equipment maintenance schedules, is a good example of a fully cooperative negotiation. Some negotiations are also characterized by the presence of a mediator, where an impartial outsider has the role of helping the parties find a compromise solution. Although often useful, this refinement offers no fundamental change to negotiation models, and we do not address it further.

There are at least two main phases to any negotiation. These are:

1. Information exchange:
  - Pre-bargining, including identification of the issues, establishing maximal limits to the issues, and agreeing on the rules
  - Issue iterations (offers and counteroffers)
2. Arriving at the outcome
  - convergence on a final contract value,
  - retention of the status quo (no change) via a walk-out by one party.

Characterizing features of negotiations have been set forth in a number of works, including [2,12,11,14,]. Of these, a key attribute is whether the negotiation involves 2 parties (bilateral) or more (multi-lateral). Multi-party negotiations have complexity that significantly exceeds that of bilateral negotiations, as players may form any of a number of different coalitions. Even in the simplest of cases, the three-party negotiation (A, B, C), one must account for any of four scenarios: no coalition, or coalition of AB,

AC, or BC. One obvious approach is to abandon negotiation altogether and utilize, for example, voting, some form of arbitration such as an auction, or perform the multi-party negotiation as a sequence of independent bilateral negotiations. Another approach, which maintains the essence of the multi-party negotiation, suggested in [15] and described further in [16], called Rubenstein's model of alternating offers, formulates the negotiation rules to explicitly disallow coalitions. Here, when one of the agents makes an offer, all other agents respond, with each agent accepting, rejecting, or cancelling (walking out). The negotiation terminates if all agents accept the offer (an agreement) or if one of them cancels. If the negotiation does not terminate, the negotiation proceeds to the next time period, another agent makes an offer, and the process repeats. Coalitions are prevented by disallowing inter-party communication. In addition, it is important to ensure that no party knows others' responses until the round is complete (otherwise parties have incentive to wait, to gain more information). The order in which parties can offer is randomized. In addition to avoiding coalitions, Rubenstein's model is also attractive because it is general; it works just as well for bilateral negotiations as it does for multi-lateral negotiations.

A second key attribute is the number of issues, i.e., the substances over which the negotiation occurs; there may be 1 (single issue) or more (multi-issue). The ability to handle multi-issue negotiations lies in the way each party evaluates an offer. Some method of normalizing among different issues is required, and expected utility provides for this. Other features of most negotiation models are whether or not they represent the influence of time, learning, strategic behaviour, and a pre-bargaining phase.

In reading the negotiation literature, it is important to recognize whether the author's perspective is *descriptive* or *prescriptive*. Descriptive models describe how negotiating parties actually behave, whereas prescriptive models prescribe how the parties should behave. For example, Gulliver in his well-known text [2] argues strongly against the use of utility theory in negotiation models because, he feels, it does not represent how people actually think (no one, he argues, actually decides based on quantification of their own and others' probabilities and preferences for various outcomes) and instead proposes two other models that capture the *cycling* and *developmental* features of negotiation, respectively. Considering Gulliver's criticism in the context of MANS, one may argue that MANS provides a degree of information accessibility and modeling power that was not available when Gulliver wrote, so that many of the complexities of how humans decide may now be effectively addressed. However true, this response neglects to recognize that MANS is not concerned with describing how humans negotiate (although it may be useful to build into MANS certain features of how humans negotiate). Rather, MANS is con-

cerned with providing humans with decision-support, i.e., prescribing how humans should decide. Therefore, the implementation mechanism is important only insofar as it provides us with desirable outcomes. In this sense, then, MANS has the same function that mathematical programming has had, except that MANS accommodates the feature of distributed decision-making that is now prevalent in the electric power industry.

#### **4.2.2 Computer-based negotiation systems**

There is a large and growing body of literature on computer-based negotiation systems, including MANS. We limit ourselves here to three important texts, published in 1994 [17], 1999 [18], and 2001 [16] that well-capture the state of the art at those times, together with a survey of some very recent literature published in a journal dedicated to the topic [19, 20, 21, 22, 23, 24] that can be conveniently found on-line at <http://journals.kluweronline.com/>.

Rosenschein and Zlotkin [17] make a strong case that “game theory is the right tool in the right place for the design of automated interactions,” arguing that despite its shortcomings in capturing human interactions, automated societies are perfectly amenable to the assumptions on which game theoretic models rest. They provide a list of attributes associated with machine interaction, including efficiency (outcomes should be Pareto Optimal), stability (no agent should have incentive to deviate from the agreed-upon available strategies), simplicity (low computational and communication requirements), distributedness (interaction rules should not require a centralized entity), and symmetry (the interaction mechanism should not arbitrarily favour one agent more than another). The work rests on the standard assumptions of game theory (rationality based on utility) together with a few more reminiscent of Rubenstein’s model, including: (a) each negotiation is independent of past or future negotiations; (b) agent-specific utility calculations may be transformed into common “system” units; (c) all functionalities (abilities) are equally accessible to all agents; (d) public agreements are binding; (e) no utility (in the form of money, for example) is explicitly transferred from one agent to another. They make the important clarification that their work is about design of machine negotiation protocols, where protocols are not about the low-level issue of how machines communicate (it is assumed that they do), but rather about a higher-level issue regarding the public rules by which machines come to agreement, such as Rubenstein’s model described above. Thus, they proceed to identify different problem domains and specify various protocols appropriate for that domain. For example, the task-oriented do-

main is one in which an agent's activity can be defined in terms of a set of tasks that it has to achieve (in contrast to domains where agents are concerned with moving its environment from one state to another, or where agents assign a worth to states and select the best state in which to move). Given a protocol, the remaining attributes necessary to characterize a negotiation are the space of possible deals, the negotiation process, and the negotiation strategy. They utilize standard game theoretic models (e.g., Zeuthen's) to analyze the influence of different strategies on outcomes.

Huhns and Stephens [18] also emphasize the importance of protocols, and they distinguish between communication protocols (e.g., KQML, KIF) and "interaction" protocols. They classify the different interaction protocols into coordination, cooperation, contract net, blackboard systems, negotiation, and market mechanisms. Their overview of each provides a useful taxonomy for more broadly understanding negotiation protocols.

Kraus [16] clearly distinguishes efforts in the area of designing agent *interaction* (i.e., coordination and cooperation) from that of designing agent *architecture*. Her efforts, relating to the former, integrates game theory with economic techniques and artificial intelligence heuristics to develop a strategic-negotiation model patterned after that of Rubenstein under assumptions similar to those of Rosenschein and Zlotkin. The importance of this work is in its detailed treatment of illustrating the generality of the proposed negotiation model in a diverse array of applications, including negotiations about data allocation, resource allocation, task distribution, pollution reduction, and hostage crises.

Reference [19] proposes a theory of negotiation developed with the intent of understanding negotiation support systems as computer-based decision systems, predicated on the idea that there are 8 different features that must be identified in order for a negotiation to be properly understood. These features are issues to be negotiated, entities involved, the acceptance region of the entities in the space of issues, the current location of the entities within that space, the strategies and movements of the entities, and negotiation rules, and the level and nature of assistance from an intervener (arbitrator or mediator). A long list of computer-based support functionalities is provided for each of the features, and a classification framework is provided in terms of the kind of entity set for which the system is used (group/peer-to-peer or organization/hierarchical) and the nature of the system's participation in the negotiation (assistance/support or autonomous negotiator). Reference [20] identifies 5 negotiation activities where mathematical modelling can provide prescriptive decision aid, and focuses on one of them, the search for agreement and improvements, in showing how it can be formalized as a MCDM gradient search problem or as a constraint proposal problem. Reference [21] identifies characterizing features

of distributive and integrative negotiations, terms first articulated by [25] to distinguish between “fixed-pie” negotiations where parties are inherently in conflict and compete over scarce resources such that when one party gets more, the other gets less (distributed) and win-win negotiations where some settlements can be better for both parties. Although integrative negotiations are generally multi-issued, they do not have to be as illustrated in the classic case where two sisters argue over an orange, one needing the juice and the other the peel. It is distributed as long as they do not know each others’ needs but immediately becomes integrative when they do. Integrative negotiations generally lead to better solutions, and the authors conclude that auctions should be considered when distributive negotiation cannot be converted to integrative, but auctions are not applicable where it is possible to for parties to learn about one another to determine opportunities and establish relationships. This theme is extended in [22] which purposes logrolling, an algorithmic method for multi-issued integrative negotiations that produces Pareto-optimal solutions through jointly improving exchange of issues such that loss in some issues is traded for gain in others resulting in overall gain for all parties. Reference [23], in focusing on computer-based business-to-business negotiations, also distinguishes between distributive and integrative negotiations and their relation to auction as in [21] and goes on to also compare norm and goal orientations in designing negotiation protocols. Reference [24] further addresses automated negotiation in the context of e-commerce applications, providing a useful negotiation taxonomy that collectively incorporates many of the attributes discussed piecemeal in the literature to date. Within this taxonomy, the parameters of the negotiation space include: cardinality (number of agents, number of issues), agent characteristics (role, rationality, knowledge, strategy, etc), environment (static or dynamic), goods (public or private), and parameters related to offers, information, and allocation. It also describes a number of proposed negotiation models and locates them in its taxonomy.

### **4.3 A multiagent negotiation system**

We have developed a MAS and instantiated agents with negotiation capabilities, resulting in a MANS. Section 4.3.1 describes the MAS implementation, and Section 4.3.2 describes the negotiation models that we have used.

### 4.3.1 MAS implementation

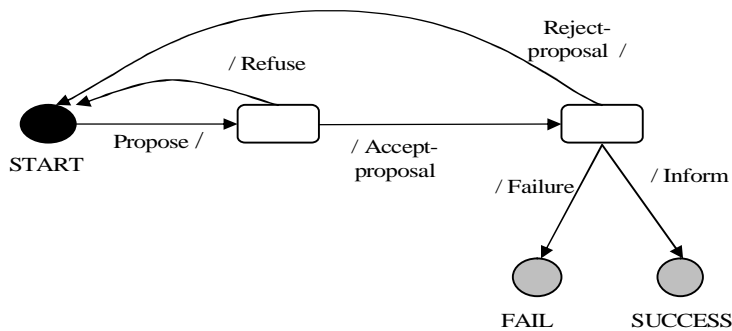
In our work, we built a platform independent Java-based API called *MASPower* [26,27,28] on top of the commercial distributed computing platform Voyager ORB [29] to instantiate agents and multiagent systems for eliciting coordinated and negotiated decision-making from power system decision-makers. Voyager supports dynamic proxy generation, naming services, synchronous and asynchronous messaging, management of multiple concurrent tasks and multiple conversation protocols, and preceptors for accessing local and remote percept sources for distributed MAS. In this framework, we used object-oriented software design methods to develop agents representing different power system entities, e.g., suppliers, transmission owners, system operators, and delivery companies. The developed software is organized into eight packages: basic and collaborative agent classes to implement agents with different functionalities, agent GUIs, tasks carried out by agents, functionalities for enabling interagent messaging, functionalities for enforcing conversation protocols, interfaces for directory services, interfaces for distributed computing inter-agent messaging, and classes for enabling interagent negotiations. Individual agents may reside on any cpu within a network as long as the cpu is running *MASPower* on top of Voyager.

We extended the federated directory service implementation of Voyager ORB to provide the ability to maintain names of currently active agents together with keywords to identify the agent's area of expertise. *MASPower* stores the directory location as an XML document, read by every newly created agent, to avoid the need to recompile a program every time the directory location is changed.

Agent communication is performed using inter-agent messaging with message interpretation being private to each agent, providing the ability to interpret the same message differently under different agent internal states. Structural elements of an inter-agent message is per FIPA-2000 recommendations [30] and include sender, receiver, content, ontology, conversation ID, protocol, reply-with, in-reply-to, and language. Multiagent conversations are managed using thread, tagged by unique conversation identifiers generated by the agent initiating the conversation. Conversation protocols were designed as finite state machines (FSM) following the COOL notations [31]. The FSM for a conversation protocol is characterized by a START state, END state, FAIL state, and a variable number of intermediate states. Transition between one state to another occurs by either sending or receiving a message with a particular performative. For example, the FIPA recommendation for the contract net protocol can be encoded as the FSM in Fig. 1. This protocol [32,33] is useful for automated

contracts in environments where all agents cooperatively work toward the same goal. The manager proposes a task, announces it, and potential contractors evaluate it (together with other announcements from other managers) and then submit bids on the tasks for which they are able to perform.

The FSM to be used by an agent depends on the role that the agent is playing in the conversation: the FSM in Fig. 1 is used by the agent responding to the initiating agent. The initiating agent uses the same FSM except that “send” and “receive” labels are interchanged for all transitions.



**Fig. 1: FSM of Contract Net Protocol in COOL Notations**


Each activity that can be undertaken by an agent in its lifetime is organized as tasks. Whenever a new task instance is created, the object registers with the agent’s task manager. A key attribute of *MASPower* is that many tasks can run concurrently within the agent.

#### 4.3.2 Negotiation protocol for individual rationality

We conceive of a society of agents organized as the electric power industry is organized, i.e., there are agents corresponding to load-serving entities, generation owners, transmission owners, and whatever centralized organizations that may exist such as ISOs, reliability authorities, and power exchanges. Decisions are made as a result of different inter-agent negotiations. Our negotiation protocol is bilateral, multi-issued (and integrative), and may be applied to decisions with or without incorporation of uncertainty. We begin by describing the protocol in terms of multilateral negotiation, without uncertainty, as it both general and simple. By avoiding the need to model uncertainty, agent decisions are made based on assessment of *utility* (or value), rather than *expected utility* (or expected

value). One attractive feature of this paradigm is that successful termination of a negotiation is guaranteed to be *pareto optimal*.

An agent that initiates the negotiation process is termed as *initiator*, and one or more responding agents are termed *responder(s)*. Our approach is patterned after that of Faratin, et al in [34] using value tradeoffs; the FSM of our protocol is illustrated in Fig. 2.



**Fig. 2: FSM of Negotiation Protocol**

The agents negotiate over values for a set of *preferentially independent* issues (pp. 101 in [6]). When two issues are preferentially independent then each agent may assert its preferences for an increasing or decreasing value of one issue without any relation to the other issue. When there are more than two issues, preferential independence is similarly defined for every subset of the set of issues and its complement. The set of issues for negotiations between any two entities usually revolve around some notion of *{quantity, quality, unit price}*. As proposed in [34], agents negotiate on the value of an issue that is within a delimited range, i.e., agents first mutually agree on a range of allowable values for every issue (pre-bargining).

In a multiagent system  $S$  consisting of at least two agents, an agent  $a$  wants to negotiate a value for a set of issues with agents in a subset of  $S - \{a\}$ ; let  $X_a = \{x_1, x_2, \dots, x_n\}$  be the set of issues about which agent  $a$  wants to negotiate, each taking values in the range specified in the set:

$$\text{range}(X_a) = \{[\min(x_1), \max(x_1)], \dots, [\min(x_n), \max(x_n)]\}$$

The set  $X_a$  is called a *negotiation set*. The agent uses a non-decreasing or non-increasing scoring function  $V(x)$  to score the value of each issue be-

tween 0 and 1<sup>2</sup>. Such functions are sufficient to model transitive preference structures. If the agent prefers an outcome  $x'$  to  $x''$  for a single issue, then  $V(x') > V(x'')$ ; if the agent is indifferent between two outcomes  $x'$  and  $x''$  then  $V(x') = V(x'')$ . *MASPower* uses a model of additive value functions (p. 90 in [6]) to get the net value of a negotiation set, through what we call the scoring function of the negotiation set, as given by:

$$V^a(X) = \sum_{i=1}^n c_i * V(x_i) \quad (1)$$

where the  $c_i$  are the relative importance of issues  $x_i$  to the agent, such that the weights are non-negative and sum to 1. Finding the relative importance of issues could be based on a subjective assessment by the human owner. This additive scoring function is the simplest multi-issued function with useful properties: two agents using additive scoring functions are sure to maximize social welfare between them<sup>3</sup>.

An agent uses either a non-decreasing or a non-increasing scoring function  $V(x_i)$  to assign a score (between 0 and 1) to a permissible value of an issue  $x_i$ . *MASPower* implements scoring functions based on the implementation in [35]: For a non-decreasing issue  $x$  defined on the domain  $[\min(x), \max(x)]$ ,


$$V(x) = (x - \min(x)) / (\max(x) - \min(x))^k \quad (2)$$

The family of curves is shown in Figure 2. The curves have an interesting interpretation in case of utility functions:  $0 < k < 1$  models risk-averse behavior. Risk-averse behavior is associated with a decision-maker who prefers an outcome with lower risk between outcomes with same expected consequences;  $k > 1$  models risk-prone behavior that is associated with a decision-maker who seeks risk; and  $k = 1$  models risk-neutral behavior (cf. pp. 148-157 in [6]). For an issue with non-increasing value function, the score obtained in eq. (2) is subtracted from one.

---

<sup>2</sup> “ $x$ ” is used to denote the negotiation issue and also the value for that issue whereas  $V(x)$  denotes the score for that value of the issue. For example, an issue  $x$  can have a value of 50 when it is defined in the range [0,100] and a score of 0.6.

<sup>3</sup> Therefore, for a negotiation process that has successfully concluded, the outcome of the negotiation is *pareto optimal* for the agents. The proof is straightforward and is available on p.164 in [11].



**Fig. 3: Value Functions Corresponding to Eq. (2)**

If the agent wants to negotiate on an issue, the agent's value of this issue is modeled solely a function of its resources, i.e., the agent uses the current state of its resources to determine a value for an issue. This is captured by the *negotiation functions*, which serve to generate a value of a single negotiation issue as a function of the agent's resources:

$$x = \sum_{i=1}^n w_i * f(r_i), \quad \sum_{i=1}^n w_i = 1 \text{ and } w_i \geq 0 \quad (3)$$

where the  $r_i$  are the "resources of the agent" to set the value for negotiation issues [34]. An agent's resources are those quantities controlled by the agent that influence the agent's decision-making with respect to one or more negotiation issues. Agent resources could include, for example, money the agent has at its disposal (where it should be clear that a *resource* of money is distinct from a negotiation issue of money). Another typical resource is available negotiation time, since an agent typically knows a particular time by which a decision is required. A resource particular to our work is equipment life. One feature of this approach, where issue values are determined by agent resources, is that issue values, at any instant during the negotiation, do not depend on the offers that this agent receives from the other agents. An important implication of this is that the agent does not behave strategically in response to bids from other agents.

Although the function  $f(r)$  can model complex dependencies, the current implementation of *MASPower* uses a family of parameterized functions for this purpose. The first step in generating the function  $f(r)$  is to model the agent's individual interpretation of the significance of  $r$ . This is

because for the same amount of resource (say time), some agents would prefer to concede rapidly when the time remaining to complete the negotiation is less, while others might prefer to hold on and not concede. If the normalized remaining resource is  $r$  (where  $r=1$  indicates all of the resource remains and  $r=0$  indicates none of the resource remains), an agent's individual interpretation of the significance of  $r$  is modeled as a function  $g(r)$  using two parameters:  $x_0$ , the normalized initial value for an issue, and  $b$ , an index (cf. [34]), via the function

$$g(r) = x_0 + (1 - x_0)r^{1/b} \quad (4)$$

The agent's negotiation function  $f(r)$  for the issue  $x$  can be written as:

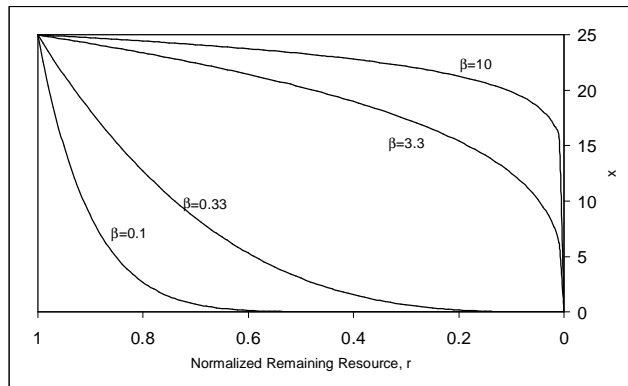
$$x = f(r) = x_{\min} + g(r) * (x_{\max} - x_{\min}) \quad (5)$$

for an issue  $x$  with increasing value function, and

$$x = f(r) = x_{\min} + (1 - g(r)) * (x_{\max} - x_{\min}) \quad (6)$$

for an issue  $x$  with decreasing value function.

Issues that depend on multiple resources are evaluated based on eq. (3). The value for an issue that is prescribed by eqs (5) and (6) always fall within the allowed interval for the issue because of the parameterization. In other words, these equations never recommend an invalid value for an issue. The influence of parameter  $b$ , which models the attitude of the agent towards a resource, is best understood from Fig. 3. When the agent has maximum resource available and as the quantity of the resource decreases (i.e.,  $r$  changes from 1 to 0), a model with  $b < 1$  concedes quickly, while there is still significant resource remaining, while a model with  $b > 1$  concedes slowly, waiting for the resource to diminish further. As the resource diminishes, the latter model concedes rapidly while the former model concedes more slowly. The family of curves with  $b > 1$  are called *Boulware* [2,11], whereas  $b < 1$  are termed *Conceder* [34].



**Fig. 4: Boulware and Conceder Negotiation Functions  $f(r)$**

The agent's decision-making algorithm under this negotiation paradigm is based on its score for the negotiation set. Suppose the agent  $a$ 's private scoring function for a negotiation set  $X$  at the time instant  $t$  is  $V^a(X)$ . This score is purely based on the internal state of this agent and is a function of its value functions, tradeoffs, current state of resources, and negotiation functions. When this agent receives an inter-agent message containing an offer  $X'$  from another agent for the same negotiation thread at a time  $t' > t$ , agent  $a$  replies with either a refuse, accept, or counterproposal, as illustrated in Fig. 2.

#### 4.3.3 Negotiation protocol – social rationality

In this section, we develop an expected utility-based model for multiagent rational decision-making, which is capable of dealing with uncertainty as well as enabling each agent to balance between its social behavior and self-interest.

As pointed out in [28], the value function based multiagent negotiation model oversimplifies the preference and attitude of the agent and precludes the possibility of modeling any uncertainty in the outcome of an action. However in practice, agents do face uncertainties in the consequences of their actions. While the Expected Utility theory has been widely adopted as the quintessential paradigm for decision-making in the face of uncertainty [8,9], the predominant view of the agent's decision-

making function has been solipsistic<sup>4</sup> in nature, based upon the principle of maximizing the individual expected utility [36] given the probability of reaching a desired state and the desirability of that state. Although this is intuitively and formally appealing, it lacks applicability in real systems consisting of multitude of interacting agents.

Given the integrity nature of power system, when designing a system in which multiple agents need to interact (coordinate) in order to achieve both *individual* (e.g. maximize individual entity's revenues) and *system* (e.g. minimize the system risk level) goals, we hold that neither egoism nor altruism are the best means to achieve globally optimal system states, but a good combination of these two aspects of interaction can yield the best global results. Because of the inherent interdependencies between the agents, an agent's decision affects not only itself, but also other agents in the multiagent system environment. It is therefore important to equip the agents within a multiagent system with a mix of self-interest and social consciousness that allows them to value the performance of the entire society over their individual performance. An agent should care about not only its individual utilities, but also all the social utilities (the utilities afforded to other agents in the entire MAS) of all possible actions while determining which action to perform. This is more reasonable when MAS is applied to the power system where entity-represented agents do have physical connections. If an agent places more emphasis on its individual utility, it is selfish in nature. On the other hand, an altruistic agent pays more attention to social utilities. A *socially rational agent* tries to maintain a balance between individual and social responsibilities [37].

Due to its intuitive and formal treatment of making decisions from a set of alternatives under uncertainty, we use the aforementioned expected utilities of the agent's actions to describe a dynamic utility calculating framework, which could provide agents a more descriptive notion of choice within a multi-agent environment. In our multiagent system, a particular agent may work in a group with a small number of agents, a loose confederation with a larger number of agents and hardly at all with the remaining agents. Thus, the calculation of social utility (includes individual utility) can be further distinguished by differentiating between the different social relationships in which an agent is engaged. We first define  $R$  as the agents' *Social Relationship Matrix* in our multiagent system:

---

<sup>4</sup> Solipsistic (unlike *Social*) agents do not explicitly model other agents within a multi-agent system.

$$R = \begin{bmatrix} r_{1,1} & r_{1,2} & \cdots & r_{1,n-1} & r_{1,n} \\ r_{2,1} & r_{2,2} & \mathbf{L} & r_{2,n-1} & r_{2,n} \\ \mathbf{M} & \mathbf{M} & \mathbf{O} & \mathbf{M} & \mathbf{M} \\ r_{n-1,1} & r_{n-1,2} & \mathbf{L} & r_{n-1,n-1} & r_{n-1,n} \\ r_{n,1} & r_{n,2} & \mathbf{L} & r_{n,n-1} & r_{n,n} \end{bmatrix} \quad (7)$$

$R$  is a  $n \times n$  matrix ( $n$  is the number of agents within the multi-agent system); each row or column corresponds to an agent. Each  $r_{i,j}$  should be a non-negative number; the value of  $r_{i,j}$  indicates agent  $i$ 's attitude toward the social relationship between itself and agent  $j$ . This means that each agent can quantitatively weight all social relationships with respect to the influence of its possible actions within the multiagent system. While different agents may have different social perspectives, the value of  $r_{i,j}$  is not necessarily equal to  $r_{j,i}$ . Therefore, in a general case the social relationship matrix  $R$  is not symmetric. This social relationship matrix could be maintained in a central secure repository<sup>5</sup>. Each agent can retrieve its own social relationship information from matrix  $R$  using its *Identification Matrix*  $I$ . This social relationship matrix coupled with individual identification matrix mechanism can be regarded as a security feature provided by the MAS platform similar to the private-public key distribution mechanism in computer system security. For instance, for the  $k$ -th agent, its identification matrix is a vector in the form of:

$$I = [i_1 \quad i_2 \quad \mathbf{K} \quad i_j \quad \mathbf{K} \quad i_n] \quad (8)$$

where  $i_j = 1$ , if and only if  $j = k$ ; otherwise  $i_j = 0$ .

Supposing the  $k$ -th agent executes a possible action  $A$ , from its perspective, the utilities afforded to each agent in the multiagent system by  $A$  can be denoted as:

$$U_A = \begin{bmatrix} SU_{k,1} \\ \mathbf{M} \\ IU_{k,k} \\ \mathbf{M} \\ SU_{k,n} \end{bmatrix} \quad (9)$$

---

<sup>5</sup> Neither the central repository nor other agents could have access to individual agent's social relationship information except that the agent is cooperative and willing to reveal its information.

where  $IU_{k,k}$  represents the individual utility to the  $k$ -th agent by that action, and all others are social utilities<sup>6</sup>. Then we can compute the expected utility for  $k$ -th agent if it carries out action  $A$  by the following equation:

$$EU(A) = I_k * R * U_A \quad (10)$$

By trying to maximize the above combined expected utility, agents can naturally take both individual and social utilities into the consideration of their decisions. The above mechanism equips the agents within multi-agent systems with a mix of self-interest and social consciousness that allows them to rationally evaluate their individual performance over the entire society. In addition, by varying the corresponding values in the relationship matrix  $R$  (if  $r_{i,j}$  is set to zero, that means agent  $i$  either neglects the social relationship ( $i \neq j$ ) between the two agents or totally removes its personal benefits from its decisions ( $i = j$ )), each agent can dynamically determine the way that it combines the individual and social utilities of all possible actions in order to make a rational decision.

## 4.4 Illustrations

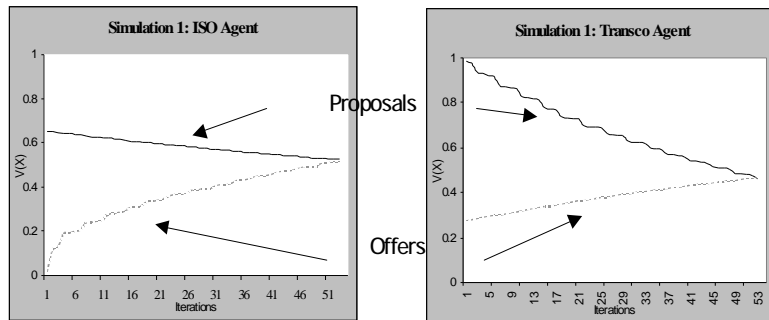
### 4.4.1 Security-related decision-making

We used the 30-bus IEEE Reliability Test System [38] under stressed operating conditions, with the following decision required: Determine by how much to operate a transmission circuit in excess of its identified rating? In the traditional vertically integrated energy industry, this decision was made by a single organization, the utility company, because it both owned and operated the transmission system. These two functions are now separated, with the ISO responsible for system operations and implementing the market based dispatch insofar as system security limits allow. However, a transmission owner is responsible for the physical integrity of the circuit, including the specification of the circuit rating, and in addition, the transmission owner receives revenues for use of the circuit in proportion to the flow. Implementations of common power system algorithms are made accessible to the agents from “C” implementations of the power system software, MATPOWER [39], and other codes. Data pertaining to the power system is maintained in a relational database to facilitate persistent

---

<sup>6</sup> In a cooperative environment, when agent  $k$  executes a possible action  $A$ , we assume that each agent would let agent  $k$  know its utility offered by action  $A$ . This is simple and enable agent to avoid modeling other agents in the system.

storage and algebraic operations on it [40]. So we have simulated a negotiation between the ISO and the Transco over the increase in circuit rating and pro-rata compensation for the transmission service. Both agents employ the value-function based negotiation model discussed in section 4.3.2. The negotiation issues are circuit rating increase and monetary compensation. The resources used by the agents are equipment life (only for the Transco), money, and negotiation time. Fig. 5 illustrates the progress of the negotiation. The negotiation concluded after 54 iterations, taking 282 seconds, when the ISO accepts an offer of 4.62 MW rating increase at \$13.13 for each MW of transmission service. A second simulation (not shown) repeated the first, except that the negotiation time resource for the Transco was decreased from 600 sec to 240 seconds, resulting in agreement after 42 iterations taking 225 seconds, at 5.54 MW rating increase at \$11.76 for each MW of transmission service. By decreasing the Transco's resource 'negotiation time,' the Transco makes coarser changes in its proposals, thereby missing several good intermediate deals.

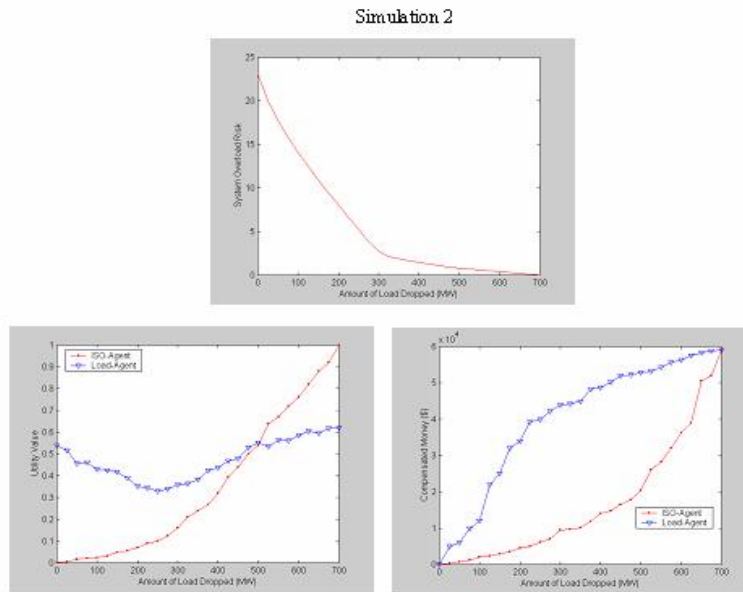


**Fig. 5: Negotiation Process Between ISO and Transco Agents**

We also did another interesting inter-agent negotiation simulation using the rational negotiation model based on expected utility described in section 4.3.3. By applying several modifications to the original system [38], we constructed a security-constrained case with system high overload risk [41]. Representing the independent system operator, the ISO-Agent is in charge of the overall system security and periodically obtains the system risk value calculated by RBSA-Agent. When it discovers high system overload risk, it immediately initiates negotiation with the Load-Agent representing the load entities at bus 13. The negotiation issues include the amount of load which the Load-Agent is willing to drop, and the amount of compensated money offered by the ISO-Agent. During the negotiation,

the Load-Agent then has to analyze the tradeoff between the compensation proposed by ISO-Agent and the expected monetary loss due to its load curtailment.

The simulation results are shown in Fig. 6. In Fig. 6, it shows that the system overload risk significantly decreases as the Load-Agent agrees to drop more and more load at bus 13, and eventually when the two agents come to an agreement that the Load-agent shields 700MW load, the system overload risk drops to zero. And the expected utilities of the two agents both mainly increase as the negotiation processes. This is because the ISO-Agent employs a negotiation function weighting system risk value more heavily than paying money, and because both agents use a socially rational negotiation model, the system security level influences the Load Agent as well as the ISO-Agent. Fig. 6 also illustrates how the compensated money evolves in the negotiation resulting in agreement between the two agents finally on \$58,974 for 700MW load dropped.




**Fig. 6: Negotiation Between ISO and Load Agents**

#### 4.4.2 Maintenance scheduling

There are thousands of high voltage power transformers in a bulk transmission system. Although power transformers have proven to be reliable in normal operation conditions with a global failure rate of 1~2% per year, the large investment in generating capacity after World War II, and continuing into the early 1970's has resulted in a large transformer population, which now is fast approaching the end of life [42]. Thus, with the large number of old electrical equipment in the power system, certainly there are numerous conflicts in scheduling equipment's maintenance activities. For the system security reason, it is very important that care should be taken in scheduling the system-wide maintenance activities to avoid severely weakening the system structure, which may lead to large system disturbance and/or customer service interrupting. We are interested in using multiagent negotiation to solve this equipment maintenance scheduling problem. Given the autonomous nature of agent, as pointed out before, inter-agent negotiation is a powerful tool to solve conflicts. Because maintenance activities would not only save money for each utility (by avoiding costly equipment failures and extend the life of electrical equipment), but also significantly improve the system reliability, the multiagent rational negotiation model developed in section 4.3.3 is suitable in this situation.

In this illustrative example, as shown in Fig. 7, each utility (transmission owner) employs several maintenance agents to be in charge of its electrical equipment's maintenance scheduling. The number of maintenance agents belonging to one utility depends on the amount of equipment the utility owns. Among the maintenance agents which belong to the same utility, they could simply rank their maintenance activities according to certain criterion, e.g. equipment failure probability, available budget etc. While in order to carry out the maintenance activity, the corresponding maintenance agent should also get approval from the ISO-Agent which is responsible for the entire system security. Using system network information, the ISO-Agent can identify allowable maintenance outage slots (number and duration) for certain time period; the requesting maintenance agents, which represent various independent organizations, negotiate among themselves to obtain more rapid approval for their maintenance activities. For example, one maintenance agent having urgent maintenance activity may be willing to negotiate with other related maintenance agents using monetary compensation in order to obtain higher priority for its maintenance activity.




**Fig. 7: Maintenance Scheduling based on MAS Negotiation**

We used the IEEE Three Area RTS-96 system [38] to perform a simulation with three maintenance agents (A, B, C), who is responsible for those power transformers in each area respectively. Based on the system expected load profile, the ISO-Agent identifies three slots allowing transformer maintenance, one at a time. Suppose the three maintenance agents all have pending maintenance tasks and certainly would like to have their maintenance work to be scheduled as soon as possible. So the three maintenance agents initiate negotiation among themselves. The negotiation issue is the amount of money one maintenance agent would be willing to compensate the others in exchange of scheduling its maintenance work first. For example, in the negotiation between agents A and B, A proposes a monetary compensation to B with the request of scheduling A's maintenance work ahead of B. If B does not satisfy with this offer, it will try to count-offer its own monetary compensation to A in exchange of scheduling B's maintenance work first instead.

During the negotiation process, each agent first evaluate the expected utility<sup>7</sup> of its opponent's offer, if this expected utility exceeds the utility of its own proposal, the agent will accept the offer, otherwise, it will try to count-offer. As seen from Fig. 8, B accepts A's offer with the agreement of 1140 \$ monetary compensation from A in exchange of scheduling A's maintenance work ahead of B's. Similarly, B accepts C's offer with the agreement of 1260 \$ monetary compensation in exchange of scheduling C's maintenance work ahead of B's; C accepts A's offer with the agreement of 1910 \$ monetary compensation in exchange of schedul-

<sup>7</sup> All the expected utilities considered in this example include both individual utility and social utility.

ing A's maintenance work ahead of C's. Thus, the overall negotiation outcome of maintenance scheduling sequence is A, C, B.



**Fig. 8 : Multiagent Negotiation among Maintenance Agents**

## 4.5 Conclusions

Multiagent system technology is now mature enough to support negotiation as a decision paradigm among different agents. As such, it is very attractive for use in addressing a fundamental difficulty inherent to operating today's power systems where we see different stakeholders simultaneously required to compete and cooperate. We have herein illustrated two representative scenarios, one on line rating increase and another on maintenance scheduling, but there are many others, including, for example, power dispatching, emergency actions, restoration, and planning. Although agent infrastructure is available for supporting negotiated decision-making, there is significant work yet to be done in evolving negotia-

tion protocols and corresponding analytical models. There is also some concern regarding the level of cyber-security necessary, and available.

We also believe that multiagent negotiation systems hold significant promise as a tool for studying different features of distributed systems such as those found within the electric power industry. For example, it appears that MANS may play a role in examining the relation between system efficiency and the level of centralization. Here, one may conceive of comparing optimization-based simulation models that compute measures of efficiency and reliability, to similar models based on MANS.

---

[1] M. Wooldridge, "Intelligent Agents," in "Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence," edited by G. Weiss, The MIT Press, Cambridge, MA., 1999.

[2] P. Gulliver, "Disputes and Negotiations," Academic Press, New York, 1979.

[3] V. Chankong and Y. Haimes, "Multiobjective Decision Making Theory and Methodology," North-Holland, 1983.

[4] D. Bernoulli, (1738), "Expositions of a New Theory of the Measurement of Risk", English translation: *Econometrica*, 1954.

[5] J. von Neumann and O. Morgenstern, "Theory of Games and Economic Behavior," John Wiley, New York, 1944.

[6] R. Keeney and H. Raiffa, "Decisions with Multiple Objectives: Preferences and Value Tradeoffs. John Wiley, New York, 1976.

[7] P. Fishburn, *The Foundations of Expected Utility*, D. Reidel Publishing Company, 1982.

[8] B. Tapan, Decision-making under uncertainty, *New York: St. Martin's Press*, 1997.

[9] P. Fishburn, Nonlinear Preference and Utility Theory, *The Johns Hopkins University Press*, 1988.

[10] R. Ackoff and M. Sasieni, "Fundamentals of Operations Research," John Wiley, New York, 1968.

[11] H. Raiffa, "The art and science of negotiation," Harvard University Press, Cambridge, Mass., 1982.

[12] J. Cross, "The economics of bargaining," Basic Books, Inc., New York, 1969.

[13] A. Strauss, "Negotiations: Varieties, Contexts, Processes, and Social Order," Jossey-Bass Publishers, San Francisco, 1978.

[14] O. Bartos, "Process and Outcome of Negotiations," Columbia University Press, New York, 1974.

- 
- [15] A. Rubenstein, "Perfect equilibrium in a bargaining model," *Econometrica* 50 (1), pp. 97-109.
- [16] S. Kraus, "Strategic Negotiation in Multiagent Environment," The MIT Press, Cambridge, MA., 2001.
- [17] J. Rosenschein and G. Zlotkin, "Rules of Encounter," The MIT Press, Cambridge, MA., 1994.
- [18] M. Huhns and L. Stephens, "Multiagent Systems and Societies of Agents," in "Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence," edited by G. Weiss, The MIT Press, Cambridge, MA., 1999.
- [19] C. Holsapple, H. Lai, and A. Whinston, "Implications of Negotiation Theory for Research and Development of Negotiation Support Systems," *Journal of Group Decision and Negotiation*, Vol 6, Issue 3 pp 255-274, 1996.
- [20] H. Ehtamo and R. Hamalainen, "Interactive Multiple-Criteria Methods for reaching Pareto Optimal Agreements in Negotiations," *Journal of Group Decision and Negotiation*, Vol. 10, pp 475-491, 2001.
- [21] G. Kersten, "Modeling Distributive and Inegrative Negotiations. Review and Revised Characterization," *Journal of Group Decision and Negotiation*, Vol. 10, pp 493-514, 2001.
- [22] M. Tajima and N. Fraser, "Logrolling Procedure for Multi-Issue Negotiation," *Journal of Group Decision and Negotiation*, Vol. 10, pp 217-235, 2001.
- [23] H. Weigand, A. dr Moor, M. Schoop, and F. Dignum, "B2B Negotiation Support: The Need for a Communication Perspective," *Journal of Group Decision and Negotiation*, Vol. 12, pp. 3-29, 2003.
- [24] A. Lomuscio, M. Wooldridge, N. Jennings, "A Classification Scheme for Negotiation in Electronic Commerce," *Journal of Group Decision and Negotiation*, Vol. 12, pp. 31-56, 2003.
- [25] R. Walton and R. McKersie, "A Behavioral Theory of Labor Negotiations," McGraw-Hill, New York, 1965.
- [26] V. Vishwanathan, V. Ganugula, J. McCalley, and V. Honavar, "A multi-agent systems approach for managing dynamic information and decisions in competitive power systems," *Proc. of the 2000 North American Power Symposium*, Oct. 2000, Waterloo, Canada.
- [27] V. Vishwanathan, J. McCalley, and V. Honavar, "A multiagent system infrastructure and negotiation framework for electric power systems," *Proc. of the IEEE Porto PowerTech Conference*, Sept., 2001, Porto, Portugal.
- [28] V. Vishwanathan, "Multiagent system applications to electric power systems: Negotiation models for automated security-economy decisions," M.S. thesis, Dept. Elec. and Comp. Eng., Iowa State Univ., Ames, 2001.
- [29] <http://www.recursionsw.com/products/products.asp>
- [30] Foundation of Intelligent Physical Agents (FIPA) specifications available at <http://www.fipa.org>.
- [31] M. Barbuceanu and M.S. Fox, "COOL: A language for describing coordination in multiagent domains," In *First International Conference on Multi-agent systems*, AAAI Press / MIT Press, 1999, pp. 17-24.

---

[32] R. Smith and R. Davis, "Framework for cooperation in distributed problem solvers," *IEEE Trans. on Systems, Man and Cybernetics* C-29(12), pp 61-70, 1980.

[33] R. Smith and R. Davis, "Negotiation as a metaphor for distributed problem solving," *Artificial Intelligence* 20, pp 60-109, 1983.

[34] P. Faratin, C. Sierra and N. Jennings, "Negotiation decision functions for autonomous agents," *Int. Journal of Robotics and Autonomous Systems*, Vol. 24:3-4, pp. 159-182, 1997.

[35] S. Talukdar and V.C. Ramesh, "A multi-agent technique for contingency constrained optimal power flows," *IEEE Transactions on Power Systems*, vol. 9-2, pp. 855-861, 1994.

[36] L. M. Hogg, Nick R. Jennings, "Socially Rational Agents – Some Preliminary Thoughts," *Proc 2nd Workshop on Practical Reasoning and Rationality*, Manchester, pp160-162, 1997.

[37] L. M. Hogg and N. R. Jennings. "Socially rational agents," *Proc. AAAI Fall symposium on Socially Intelligent Agents*, pages 61--63. AAAI, November 1997.

[38] IEEE reliability test system task force of the application of probability methods subcommittee (1999), "The IEEE reliability test system – 1996", *IEEE Transactions on Power Systems*, Vol.14, No. 3, 1999, p1010-1018.

[39] Power Systems Engineering Research Center. MATPOWER, version 2.0. available at <http://www.pserc.cornell.edu/matpower/>.

[40] Lutriss Technologies. InstantDB, version 3.26, [instantdb.enhydra.org](http://instantdb.enhydra.org).

[41] Z. Zhang, V. Vishwanathan, and J. McCalley, "A Multiagent Security-Economy Decision Support Infrastructure for Deregulated Electric Power Systems," *Proc., 7<sup>th</sup> International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*, Naples, Italy, Sep. 2002.

[42] B. Sparling, "Transformer Monitoring: Moving Forward from Monitoring to Diagnostics", *Transmission and Distribution Conference and Exposition, 2001 IEEE/PES*, Volume: 2, 2001, Page(s): 960 -963 vol.2.

Field Code Changed