

Satisficing Negotiations

James K. Archibald, *Senior Member, IEEE*, Jared C. Hill, *Student Member, IEEE*, F. Ryan Johnson, and Wynn C. Stirling

Abstract—Negotiation procedures that are founded on the doctrine of individual rationality, where each participant is committed to maximizing its own satisfaction, are limited in their ability to accommodate the interests of others, and therefore, may unnecessarily constrain the negotiability of a decision maker, particularly in cooperative environments. Satisficing game theory provides a distinct alternative to the hyperrationality of conventional rational choice by waiving reliance on the individual rationality premise and offering an approach to negotiatory decision making that is based on a well-defined mathematical notion of satisficing, or being good enough, that permits the modeling of complex interrelationships between agents. This approach provides a mechanism to compute the attitude, or degree of conflict or contentedness, of the negotiators. Examples illustrate both single-round and multiround satisficing negotiation protocols.

Index Terms—Altruism, Bayesian networks, game theory, negotiations, rationality, satisficing, sociology.

I. INTRODUCTION

NEGOTIATION is an iterative decision-making process between independent decision-making entities as they attempt to reach a joint decision that is acceptable to all participants. Game theory, as developed by von Neumann and Morgenstern [1], provides a mathematical framework within which multi-agent decision problems can be represented and by which negotiatory processes can be evaluated. The standard application of this theory requires each player to form a utility function that quantifies the benefit that accrues to it as a consequence of the actions that it and all other players may take. A *strategic-form game* is created by juxtaposing the utilities of all players into an array that is available to all players. Each player then may assess the opportunities for cooperation and conflict. A *solution concept* is a rule that defines what it means for a decision vector to be acceptable to all players in the light of the conflict/cooperation environment. A *negotiation* is a solution concept whereby the players iteratively modify their proposed solutions in an attempt to reach a mutually acceptable decision vector. Such an agreement, if reached, is called a *compromise*. One way in which negotiation differs from other solution concepts such as Nash equilibria is that a compromise is not guaranteed to exist. If one is not achieved, then the negotiation defaults to an impasse. To avoid vacuous situations, we assume that, once a compromise is reached, all participants will enact the negotiated solution.

Many negotiation concepts have been devised within the general game-theoretic framework, including voting, auctions, bar-

Manuscript received January 29, 2004; revised July 12, 2004.
The authors are with the Electrical and Computing Engineering Department, Brigham Young University, Provo, UT 84602 USA (e-mail: jka@ee.byu.edu; hilljc@byu.edu; ryanjohnson@ece.cmu.edu; wynn@ee.byu.edu).
Digital Object Identifier 10.1109/TSMCC.2006.860571

TABLE I
PAYOFF MATRIX IN ORDINAL FORM FOR (a) THE PRISONER'S DILEMMA AND (b) THE BATTLE OF THE SEXES GAME

P_1	P_2		H	S	
	C	D		D	B
C	(3, 3)	(1, 4)	D	(4, 3)	(2, 2)
D	(4, 1)	(2, 2)	B	(1, 1)	(3, 4)

(a)

(b)

Key: 4 = best; 3 = next best; 2 = next worst; 1 = worst

gaining, arbitration, and market equilibrium mechanisms [2], [3]. While these concepts are all very different from each other and each has its own natural application domain, they share one important feature: by virtue of being couched under the rubric of game theory, they are all consistent with the classical economics hypothesis advanced by Bergson and Samuelson, which asserts that individual interests are fundamental, i.e., that social welfare is a function of individual welfare [4], [5]. This is the doctrine of *individual rationality*, that is, each individual should act in its own self-interest. Under this doctrine, all other agents are viewed, to varying degrees depending on the context, as competitors that act to constrain one's actions. The injunction to each agent is very simple—it should optimize (unconstrained if possible, constrained if necessary). Individual rationality is certainly an appropriate paradigm for decision makers who are primarily competitive. Negotiators operating in such scenarios would be inclined to use various posturing devices such as exaggerations, threats, and even outright deception. Much of negotiation theory is designed to deal with such tendencies.

Many negotiation scenarios, however, possess a stronger cooperative than competitive nature, and in such cases, notions of collaboration, compromise, and altruism are more applicable descriptors than competition, exploitation, and avarice. To illustrate the difference, let us compare two well-known games that are often used as prototypes of social situations: the Prisoner's Dilemma (PD) and the Battle of the Sexes (BOS). Payoff arrays in ordinal form for these two games are illustrated in Table I.

PD is a ubiquitous model of economic behavior, and has been applied to a wide variety of situations ranging from arms races to oligopoly pricing. The most common instantiation of the game involves two players (P_1 and P_2) who may either cooperate (C) or defect (D). If one player cooperates and the other defects, the one who defects receives the best payoff while the one who cooperates receives the worst payoff. If both defect, they both receive the next-to-worst payoff, and if both cooperate, they both receive the next-to-best payoff. The game has one Nash equilibrium, (D, D), resulting in the next-to-worst outcome for both.

BOS involves a man (H) and a woman (S) who plan to meet in town for a social function. She prefers to go to the ballet

(B), while he prefers the dog races (D). Each also prefers to be with the other, however, regardless of venue. This game is a prototype for economic scenarios where decision makers (e.g., firms) are trying to choose between competing standards. Although each firm has its own preferences, both firms would sell more products if they were to adopt a common standard. Resorting to such tactics as threats, bluffs, and deception is likely only to exacerbate the problem. To avoid an impasse, it will be necessary for some form of compromise to occur, and the question is, what could be the basis for cooperation? Classical game theory offers two Nash equilibria, namely, (D, D) and (B, B), but does not provide a unique resolution, hence is completely ineffective in resolving the impasse.

PD is a game of mixed motive and serves as an appropriate model of behavior when the opportunity for exploitation exists. Cooperation, though possible, incurs great risk, while defection, even though it offers diminished rewards, protects the participant from catastrophe. By contrast, BOS is largely a game of coordination, in which the interests of the agents largely coincide. With this game, there is little (though not zero) opportunity for exploitation; either both players win (although not to the same degree) or both lose. While individual rationality may be an appropriate paradigm with which to analyze PD, that paradigm loses much of its instrumentality as an effective way to analyze BOS. Both games have a significant opportunity for cooperation, but failure to do so, while detrimental to PD, can be catastrophic for BOS. These games also differ in their amenability to negotiation. Negotiation with PD is conceptually trivial, since both players receive a greater reward by cooperating than by defecting. Negotiation with BOS, however, is not so straightforward, since the rewards for cooperating are unequal. It seems that more must be taken into account when negotiating than the payoff array.

One concept that must be taken into account is altruism. An altruistic agent takes into consideration the preferences of others when defining its preferences (benevolently or malevolently—altruism is deontologically neutral—Taylor [6] calls this positive and negative altruism). An agent is *categorically altruistic* if it relinquishes its demand to optimize its own benefit, without regard for others, in all circumstances. A classical way to do this is to form one's utility as a linear combination of the utilities of all players [6]. The player then proceeds to invoke the usual solution concept, such as Nash equilibrium. The key feature of this approach is that the player's preferences are unconditional. It has irrevocably re-defined its preferences in a way that obviates its narrow self-interest. For example, consider the BOS game. If both H and S were categorically altruistic, they would each prefer to go to the venue that is more appealing to the other. Unfortunately, the results would be the worst possible outcome for both.

An alternative to the categorical approach is *situational altruism*, where the player conditionally relinquishes its narrow self-interest if, but only if, the other wishes to take advantage of the offered largesse. Otherwise, the agent would be governed by its egoistic preferences and would avoid needless sacrifice. In the BOS game, suppose that H is not a stereotypical machoistic male who has little consideration for the feelings of his partner. Although he prefers dog races to ballet, let us cast

him as a somewhat sensitive fellow who wants his friend to enjoy herself. He feels this way strongly enough to be willing to moderate his preference for the dog races if, but only if, S really hates that environment. He may express this attitude by defining two utility functions, one under the assumption that S detests dog races, and the other under the assumption that she tolerates them. Such preferences are conditional for H , in that he does not commit to either preference independently of S 's attitude. These utilities can be defined without H even knowing S 's attitude about dog races. Notice, also, that it is possible for H to make these conditional evaluations without knowledge of S 's attitude about ballet.

Conditional preferences are difficult to express via the utilities that are conventionally used in game theory. The reason for this difficulty is that, with von Neumann–Morgenstern game theory, each agent defines its preferences as a function of the possible *actions* of itself and other players, and orders these preferences by considering only the effect that the outcomes have on its own level of satisfaction. It is not until the payoffs are juxtaposed in the payoff array that opportunities for conflict or cooperation are revealed. If, on the other hand, preferences were formed as functions of other players' *preferences* as well as one's narrow self-interest, then the players would possess the capability of forming their preferences selfishly, benevolently, malevolently, or indifferently with respect to the preferences of others.

Optimization (constrained or unconstrained) is a seemingly incontrovertible solution concept, especially in a single-agent context. In a multi-agent context, however, optimization is not always a well-defined concept. On the one hand, the players can adopt a "bottom-up" approach, where each optimizes its own behavior. However, such behavior may not generate optimal group performance. As Arrow's impossibility theorem establishes, it is generally not possible simultaneously to optimize both individual and group performance [7]. Alternatively, the group could adopt a "top-down" approach, whereby, acting as a single entity, each player's actions are specified so as to optimize group behavior. So doing, however, may require unacceptable degradations of individual performance and will not generally result in an enforceable solution concept. Essentially, optimization is an individual activity. As observed by Luce and Raiffa, "the notion of group rationality is neither a postulate of the model nor does it appear to follow as a logical consequence of individual rationality" ([8], p. 193).

This paper presents a theory of negotiation that is *not* based on the doctrine of individual rationality. Rather, it is based on a mathematically precise notion of being "good enough" that is fundamentally different from, and not an approximation to, being "best." Key attributes of this approach include the following:

- it naturally accommodates sophisticated social behaviors such as cooperation, compromise, and negotiation;
- it does not depend on optimization as the criterion for defining quality; alternative criteria for quality are introduced;
- it leads to well-defined solutions procedures and is amenable to a precise mathematical characterization.

II. SATISFICING GAMES

The term “satisficing” was first appropriated by Simon [9], [10] as a species of “bounded rationality.” Under his usage, a decision maker searches for the optimal solution but terminates the search when an option is deemed to be good enough. The concept of being “good enough,” in this context, is that the performance attributed to the option met or exceeded the decision maker’s heuristic “aspiration level.” This usage takes into consideration exigencies of practical decision making such as the informational and computational constraints that exist in real-world situations. It is similar in philosophy to individual rationality since, if a decision maker could optimize, it surely should do so. Only the real-world constraints on its capabilities prevent it from achieving the optimum, but the *ideal* of optimality remains intact.

For single-agent low-dimensional problems, it may be straightforward to specify the aspiration levels. But, with multiple-agent systems, it may be difficult to define a notion of group aspirations, and perhaps even more difficult to reconcile group aspirations with individual aspirations. Even if this were possible, interdependence between decision makers can be complex and aspiration levels can be conditional (what is good enough for me may depend upon what is good enough for you). It seems that Simon’s notion of satisficing cannot be easily extended in a systematic way to the multi-agent case.

This paper also employs satisficing to mean good enough, but not in the sense of bounded rationality. There are two significant differences between satisficing *à la* Simon and satisficing as used here.

- 1) Simon’s usage is an approximation to being best (constrained from achieving that ideal by practical limitations). By contrast, *satisficing as used here treats being good enough as the ideal* (rather than an approximation). Thus, satisficing, as employed in this paper, is *not* a species of bounded rationality, in the sense of being an approximation to being optimal.
- 2) Simon’s standard for being good enough is *extrinsic*, that is, as with optimization, options are evaluated with respect to attributes that are not part of the option. In the case of optimization, evaluations are made relative to other options. In the case of aspiration levels, evaluations are made relative to an externally supplied aspiration level. By contrast, satisficing as defined here involves *intrinsic* evaluations, that is, intra-option evaluations of multiple attributes of each option without reference to sources of information outside the option.

One way to form intrinsic comparisons is to form dichotomies, that is, to define two distinct sets of attributes for each option and either to reject or fail to reject the option on the basis of comparing these attributes. Such dichotomous comparisons are intrinsic since they do not reference anything not directly relating to the option. Dichotomies are the fundamental building blocks of everyday personal choices. Attached to virtually every nontrivial option are attributes that are desirable and attributes that are not desirable. People are naturally wont to evaluate the upside versus the downside, the pros versus the

cons, the pluses versus the minuses, the benefits versus the costs. One simply evaluates tradeoffs option by option—putting the gains and the losses on the balance to see which way it tips. The result of evaluating dichotomies in this way is that the benefits must be at least as great as the costs. In this sense, such evaluations provide a distinct notion of being good enough.

By separating the positive (benefit) and negative (cost) attributes of an option, we explicitly raise the issue of commensurability. However, this issue is also implicitly present with conventional utilities, since they also typically involve both benefits and costs, and the decision maker must somehow determine the relative significance of these attributes. A typical conventional procedure is to form a linear combination of the positive and negative attributes, with the weighting coefficients being chosen to correspond to significance. The distinction between the traditional approach and our approach is that we do not aggregate the different attributes into a single function, but instead keep them separate. At the end of the day, both approaches require the subjective evaluation of significance by the designer, who must formulate some rational notion of commensurability by appropriating or inventing a system of units. The issue was put succinctly by Hardin: “Comparing one good with another is, we usually say, impossible because goods are incommensurable. Incommensurables cannot be compared. Theoretically, this may be true; but in real life incommensurables *are* commensurable. Only a criterion of judgment and a system of weighing are needed” ([11], emphasis in original). To make dichotomous comparisons meaningful, we must express the benefit and cost in the same units by insisting that they be normalized (as will be developed subsequently).

We define an option as being *satisficingly rational* if the gains obtained by adopting it are at least as great as the losses so incurred. This notion of rationality is weaker than individual rationality (which involves extrinsic comparisons); it provides an explicit definition of what it means to be “good enough.” Whereas individual rationality may be characterized as an attitude of “nothing but the best will do,” satisficing rationality may be characterized as an attitude of “getting what you pay for.” Individual rationality is rigid and demanding; satisficing rationality is ameliorative and flexible. There can be only one individually rational option (or an equivalence class of them) for a given optimality criterion, but there can be several satisficingly rational options for a given satisficing criterion. It is easy to show, however, that, if computed by the same performance criteria, an optimal decision will also be satisficing. Both notions fit Nozick’s definition of *instrumental rationality* as “the effective and efficient pursuit of given goals” [12]. Furthermore, as Arrow observed: “Among the classical economists, such as Smith and Ricardo, rationality had the limited meaning of preferring more to less” [13]. Individual rationality has taken this rather primitive injunction to its extreme instantiation as optimization, but that does not imply the impossibility of other notions of rational behavior. As we will show, satisficing is designed to permit the agents to extend their spheres of interest beyond the self, thereby facilitating negotiation. We introduce the mathematical concepts first for the individual, then extend to the multi-agent case.

A. Single-Agent Satisficing

Definition 1: Let U denote a finite set of options. A *mass function* p is a mapping of U onto the unit interval such that

$$p(u) \geq 0 \forall u \in U, \quad \sum_{u \in U} p(u) = 1.$$

A common use of mass functions is to characterize the probability distribution of a discrete random variable. In our application, however, we employ mass functions in a way analogous to probability theory, but with a different interpretation. In the interest of making intrinsic comparisons of attributes of options, we will define two mass functions.

Definition 2: A *selectability mass function* p_S is a mapping that characterizes the degree of support that is attributed to $u \in U$ in the interest of accomplishing whatever fundamental goal is relevant to the decision maker.

A *rejectability mass function* p_R is a mapping that characterizes the degree to which $u \in U$ consumes whatever resources are at the disposal of the decision maker as it takes action.

Resources may consist of such things as money, fuel, exposure to hazard, or any other undesirable consequences that are distinct from achieving the fundamental goal of taking action. The tradeoff between achieving success and consuming resources identifies the satisficing options.

Definition 3: An option u is said to be *satisficing at negotiation level q* if the degree of support for its implementation is at least as great as q times the degree to which it consumes resources. The *satisficing set at negotiation level q* is

$$\Sigma_q = \{u \in U : p_S(u) \geq qp_R(u)\}. \quad (1)$$

Nominally, we will take $q = 1$, but it may be adjusted in the course of negotiation as means of lowering standards in an attempt to reach a compromise in multi-agent applications.

B. Multiple-Agent Satisficing

The advantage of using mass functions as utilities is that they can be extended to the multiple-agent case analogous to the way probability mass functions can be extended to the multivariate case to characterize the distribution of multiple random variables.

Definition 4: Consider a set of N agents, and let U_i denote the option set for the i th agent. The set of joint options is the product set $\mathbf{U} = U_1 \times U_2 \times \dots \times U_N$, and the elements of \mathbf{U} are vectors of the form $\mathbf{u} = (u_1, u_2, \dots, u_N)$, where $u_i \in U_i$. The *joint selectability mass function* $p_{S_1 S_2, \dots, S_N}$ is a mapping from \mathbf{U} to the unit interval such that

$$p_{S_1 \dots S_N}(u_1, \dots, u_N) \geq 0 \forall (u_1, \dots, u_N) \in \mathbf{U},$$

and

$$\sum_{\substack{u_i \in U_i \\ i=1, \dots, N}} p_{S_1 \dots S_N}(u_1, \dots, u_N) = 1.$$

The *joint rejectability mass function* p_{R_1, \dots, R_N} is defined similarly.

Analogous to the way univariate probability mass functions are obtained as marginals of joint probability mass functions, we may extract individual decision maker selectability and rejectability mass functions as marginals; namely

$$p_{S_i}(u_i) = \sum_{\substack{u_j \in U_j : j \neq i \\ j=1, \dots, N}} p_{S_1 \dots S_N}(u_1, \dots, u_N). \quad (2)$$

The rejectability marginal p_{R_i} for decision maker i is defined similarly.

Definition 5: A *satisficing game* [14]–[24] is the triple $\{\mathbf{U}, p_{S_1, \dots, S_N}, p_{R_1, \dots, R_N}\}$. The *jointly satisficing solution* is the subset of all option vectors such that the joint selectability is at least as great as the negotiation index times the joint rejectability, that is

$$\Sigma_q = \{(u_1, \dots, u_N) \in \mathbf{U} : p_{S_1 \dots S_N}(u_1, \dots, u_N) \geq qp_{R_1 \dots R_N}(u_1, \dots, u_N)\}. \quad (3)$$

The individually satisficing solutions for each agent are obtained from the marginal selectability and rejectability mass functions, yielding the *individually satisficing solutions*

$$\Sigma_q^i = \{u_i \in U_i : p_{S_i}(u_i) \geq qp_{R_i}(u_i)\}. \quad (4)$$

The *satisficing rectangle* is the product set of the individually satisficing sets, namely

$$\mathfrak{R} = \Sigma_q^1 \times \dots \times \Sigma_q^N. \quad (5)$$

The jointly satisficing set Σ_q represents the subset of option vectors that are collectively satisficing for the group, in the sense that the benefits to the group dominate the costs to the group. However, it is important to appreciate that this concept does *not* presuppose that there is a cohesive notion of group preference. If the group is purely competitive, such as would be the case with a zero-sum game, then the group “preference” may be to oppose each other, and the individual preferences as obtained as marginals will be consistent with narrow self-interest. On the other hand, if the group is committed to achieving a coherent collective goal, then a well-defined group preference may obtain, and the individual preference marginals will be consistent with cooperative behavior, even at the expense of individual benefit.

In general, the satisficing rectangle will not be the same as the jointly satisficing set; they may even be disjoint. However, the following theorem relates the two sets.

Theorem 1: The Negotiation Theorem: If u_i is individually satisficing for agent i , that is, $u_i \in \Sigma_q^i$, then it must be the i th element of some jointly satisficing vector $\mathbf{u} \in \Sigma_q$.

Proof: We will establish the contrapositive, namely, that if u_i is not the i th element of any $\mathbf{u} \in \Sigma_q$, then $u_i \notin \Sigma_q^i$. Without loss of generality, let $i = 1$. By hypothesis, $p_{S_1, \dots, S_N}(u_1, \mathbf{v}) < qp_{R_1, \dots, R_N}(u_1, \mathbf{v})$ for all $\mathbf{v} \in U_2 \times \dots \times U_N$, so $p_{S_1}(u_1) = \sum_{\mathbf{v}} p_{S_1 \dots S_N}(u_1, \mathbf{v}) < q \sum_{\mathbf{v}} p_{R_1, \dots, R_N}(u_1, \mathbf{v}) = qp_{R_1}(u_1)$, hence $u_1 \notin \Sigma_q^1$. ■

Thus, if an option vector is individually satisficing, it is part of a jointly satisficing vector, although it need not be part of all

jointly satisficing vectors. The converse, however, is not true: if u_i is the i th element of a jointly satisficing vector, it is not necessarily individually satisficing for agent i . The significance of this theorem for negotiation is that *no one is ever completely frozen out of a deal—every decision maker has, from its own perspective, a seat at the negotiating table*. This is perhaps the weakest condition under which negotiations are possible.

Definition 6: The *compromise set at negotiation level q* is the intersection of the jointly satisficing set and the satisficing rectangle: $\mathcal{C} = \Sigma_q \cap \mathcal{R}$.

The elements of \mathcal{C} comprise the set of options that are simultaneously good enough for the group and good enough for each individual. If this set is empty, then successive rounds of negotiation may be performed by incrementally lowering the negotiation index q until $\mathcal{C} \neq \emptyset$.

C. Interdependence

Since the behavior of the group is dependent on the structure of the joint selectability and rejectability mass functions, it is imperative that we understand exactly how these functions are created. To understand this process, it is necessary to define a more fundamental concept, that of *interdependence*. An act by any individual member of a multi-agent system has possible ramifications for the entire system. Some participants may be benefited by the act, some may be damaged, and some may be indifferent. Furthermore, although an individual may perform the act in its own interest or for the benefit of others or the entire system, the act is usually not implemented free of cost. Resources are expended, or risk is taken, or some other penalty or unpleasant consequence is incurred, perhaps by the individual whose act it is, perhaps by other participants, and perhaps by the entire system. Although these undesirable consequences may be defined independently from the benefits, the measures associated with benefits and costs cannot be specified independently of each other, due to the possibility of interaction. A critical aspect of modeling the behavior of such a system, therefore, is the means of representing the interdependence of both positive and negative consequences of all possible joint options that could be undertaken.

Definition 7: The *interdependence mass function* $p_{S_1, \dots, S_N R_1, \dots, R_N}$ is a mapping from $\mathbf{U} \times \mathbf{U}$ to the unit interval such that

$$p_{S_1, \dots, S_N R_1, \dots, R_N}(u_1, \dots, u_N; v_1, \dots, v_N) \geq 0$$

and

$$\sum_{\substack{u_i \in U_i, v_j \in U_j \\ i, j = 1, \dots, N}} p_{S_1, \dots, S_N R_1, \dots, R_N}(u_1, \dots, u_N; v_1, \dots, v_N) = 1.$$

The joint selectability and rejectability mass functions may then be obtained from the interdependence function as

$$p_{S_1, \dots, S_N}(u_1, \dots, u_N) = \sum_{\substack{v_i \in U_i \\ i=1, \dots, N}} p_{S_1, \dots, S_N R_1, \dots, R_N}(u_1, \dots, u_N; v_1, \dots, v_N) \quad (6)$$

$$p_{R_1, \dots, R_N}(v_1, \dots, v_N) = \sum_{\substack{u_i \in U_i \\ i=1, \dots, N}} p_{S_1, \dots, S_N R_1, \dots, R_N}(u_1, \dots, u_N; v_1, \dots, v_N). \quad (7)$$

A useful way to view the interdependence function is that each decision maker possesses two selves, or roles. One self considers only the positive, or selectable, attributes of the options under consideration, and the other self considers only the negative, or rejectable, attributes. The interdependence function then describes the collective attitude of the group when considering both selves of every member of the group with respect to selecting option vector (u_1, \dots, u_N) and rejecting option vector (v_1, \dots, v_N) . This structure provides a framework within which all conceivable relationships can be expressed. The special case when $v_i = u_i$ characterizes the conflict that arises between the two selves because of the desire to select an option on the basis of its expediency, but also desiring to reject it because of its expense. In the single-agent case, it is reasonable to assume that the criteria that define selectable attributes are distinct from the criteria that define rejectable attributes. In that case, because selectability and rejectability are independent, the interdependence function would factor into the product $p_{SR}(u; v) = p_S(u)p_R(v)$, but this factorization is not required by the theory. In the multiple-agent case, however, one participant's rejectability, say, may influence another player's selectability, so it is *not* generally true that the interdependence function factors into the product of the joint selectability and rejectability functions.

III. A SOCIOLOGY FOR NEGOTIATION

For cooperative negotiatory scenarios where the agents must work together to achieve a group goal, it is imperative that they function according to a sociology that supports their cooperative requirements. With the satisficing approach, the interdependence mass function characterizes all of the interconnections between the participants. They may be derived either from the perspective of individual rationality, or they may be derived from the perspective of coordination and collaboration. The interdependence function is a mathematical encoding of the sociology of the system. Its construction is the most critical aspect of satisficing game theory. The reader may have already noticed that the structure of a satisficing game is reminiscent of a Bayesian network. Recall that a Bayesian network is a directed acyclic graph (DAG) consisting of nodes and edges, where the nodes represent the variables of the system, and the edges represent a conditional probabilistic description of how the instantiation of a parent node at a particular value influences a child node.

Similar to the way a Bayesian network is defined, we may also employ the tools of graph theory to express a multi-agent system as a DAG. To distinguish between the probabilistic application and our context, we will refer to such networks as *praxeic networks*. (Praxeology is the science of efficient action). A praxeic network for an N -agent system consists of $2N$ nodes, with each participant having two nodes associated with it—one for its selectability self and one for its rejectability self. The




Fig. 1. Praxeic network for a three-agent system.

variables associated with these nodes are the options available to the decision maker and the edges represent the influence that one agent's self has on another agent's self. These linkages consist of conditional selectability or conditional rejectability functions.

Each application will have a different network and a corresponding independence function that is unique to that network. As an example, consider the graph displayed in Fig. 1. This graph corresponds to a three-agent system where the selectability of agent 1 influences the selectabilities of agents 2 and 3 and the rejectability of agent 2, and the selectability of agent 2 influences the selectability of agent 3 and rejectabilities of agents 1 and 3. The corresponding interdependence function is

$$\begin{aligned}
 & p_{S_1 S_2 S_3 R_1 R_2 R_3}(u_1, u_2, u_3; v_1, v_2, v_3) \\
 &= p_{S_3 | S_1 S_2}(u_3 | u_1, u_2) p_{R_1 | S_2}(v_1 | u_2) p_{R_2 | S_1}(v_2 | u_1) \\
 & \quad p_{R_3 | S_2}(v_3 | u_2) p_{S_2 | S_1}(u_2 | u_1) p_{S_1}(u_1). \quad (8)
 \end{aligned}$$

$p_{S_3 | S_1 S_2}(u_3 | u_1, u_2)$ is a conditional selectability mass function characterizing the degree of support that agent 3 places on selecting option u_3 given that agent 1 were to prefer to select option u_1 and agent 2 were to prefer to select option u_2 . Similarly, $p_{R_1 | S_2}(v_1 | u_2)$ represents the conditional rejectability that agent 1 places on option v_1 given that agent 2 were to prefer to select option u_2 . The other terms in this expression may be interpreted in like manner.

There are two major distinctions between Bayesian networks and praxeic networks. First, whereas Bayesian networks deal with the epistemological problem of what to believe, praxeic networks deal with the problem of how to act. Second, whereas Bayesian networks deal with one notion of probability—a measure of the degree of belief—praxeic networks deal with two “probability-like” notions: selectability, which characterizes the degree of suitability of options, and rejectability, which characterizes the degree of resistibility of options. These functions possess the same mathematical structure as do probability mass functions, with the two most important properties being the concepts of conditioning and independence.

A third distinction between Bayesian and praxeic networks is often present; Bayesian networks define connections between propositions, whereas praxeic networks define connections between agent selves. With Bayesian networks, the vertices represent interacting elements (e.g., the states of Nature), and the edges represent the flows of belief influence between them. The vertices of a praxeic network, however, represent the sets of

actions that are possible for the agent selves, and the edges characterize the flows of action influence between the selves.

A key feature of Bayesian networks is that the joint probability mass function is constructed from conditional relationships via the chain rule of probability theory. Praxeic networks share this same computational feature. It is a well known aspect of probability theory that it is often much easier to compose a joint distribution from conditional distributions by means of the chain rule than to stipulate it directly. The conditional mass functions represent hypothetical situations that are often quite simple to evaluate. They may be viewed as production rules of an expert system. The power of this approach is that the marginals can be computed via Pearl's Belief Propagation Algorithm [25], thereby establishing the degree of belief support for each state of nature. (Although the general fully linked problem is *NP*-hard, fortunately, many interesting systems are only sparsely linked, resulting in a greatly reduced computational burden.)

Praxeic networks possess an analogous interpretation. The conditional selectability and rejectability mass functions represent hypothetical situations that may occur in a society, and correspond to behavioral rules that define how the members of the society function. For example, given the hypothetical “agent i selects option u_i ,” then agent j should condition its rejectability of option v_j as $p_{R_j | S_i}(v_j | u_i)$. Using such hypotheticals, the interdependence function can be constructed by the chain rule, and Pearl's algorithm may then be applied to obtain the joint and marginal selectability and rejectability mass functions, thereby solving the satisficing game.

By exploiting and adapting the mathematical structure of mass functions, we may define a systematic design methodology for the synthesis of a multi-agent system that incorporates negotiation.

- 1) Form operational definitions of selectability and rejectability, and represent each agent's selectability and rejectability selves as nodes.
- 2) Define the influence flows between agent selves, that is, the ways in which each agent's preferences influence the preferences of others. Represent these influence flows as directed edges linking the appropriate nodes. If there is not a *direct* influence between two nodes, then no edge connects them.
- 3) Associate a conditional selectability or conditional rejectability mass function (as appropriate) with each edge. These functions represent the degree of influence that exists between the agent selves.
- 4) Compute the marginal selectability and rejectability functions for each node and the joint selectability and joint rejectability functions for the entire group (or subgroups if appropriate) using Pearl's Belief Propagation Algorithm [25] or the sum-product rule of factor graphs [26].
- 5) Compute the joint and individually satisficing sets for each node.
- 6) If the satisficing rectangle and the jointly satisficing set are disjoint, then enter into negotiations by incrementally lowering the negotiation index and repeating the previous step until the intersection is nonempty. The resulting options are then both satisficing for the group and for each

individual. If the intersection contains more than one element, the option that maximizes the group benefit is a logical one to be chosen.

IV. ATTITUDE

An important aspect of negotiation is the sense that the agents are content or conflicted by whatever settlement is achieved. Unless a negotiated decision is optimal, in some sense, for all participants, the possibility for discontent may exist. A consequence of the individual rationality paradigm is that the decision-maker dispassionately does what should be done under that behavioral regime. On the other hand, replacing individual rationality and its attendant demand for optimization with satisficing methodology provides an opportunity for the decision makers to form assessments of the quality of proposed compromises. In particular, it admits a measure of attitude, or disposition, of the decision makers that is not subjective and, though it admits anthropomorphic metaphors, does not rely upon them for its validity.

It is fortunate if an option that conserves resources (low rejectability) also achieves the goal (high selectability). If such an option is available, the decision maker would be in a state of contentment with respect to that option. Many interesting decision problems, however, are such that the only possible options that achieve the goal are relatively expensive, hazardous, or have other undesirable side effects. A decision maker in this situation is in a state of conflict. Fortunately, satisficing theory provides a natural mathematical method to define such notions as contentment and conflict. Since it is based upon the mathematics of probability theory, we may categorize the elements of the satisficing set with respect to these dispositions by adapting the notion of entropy to a praxeological context. The entropy of a mass function p is given by $H(p) = -\sum_{u \in U} p(u) \log_2 p(u)$. The classical Shannon information-theoretic interpretation of entropy is a measure of the uncertainty associated with a probability mass function. For example, if p is a probability mass function and $p(u) \approx 1$, then $-\log_2 p(u) \approx 0$, which confirms that, since the occurrence of the event u is highly probable, there is little uncertainty associated with its occurrence (or, equivalently, observing that u occurred does not greatly reduce uncertainty). Conversely, suppose $p(u') \approx 0$. Then $-\log_2 p(u')$ is large, indicating that great uncertainty is associated with predicting the occurrence of that event, or, equivalently, uncertainty in the outcome is greatly reduced by the occurrence of u' . Entropy is the average value of uncertainty over all $u \in U$, and admits two interpretations. On the one hand, $H(p)$ is a measure of the average uncertainty in the outcome of an experiment governed by p before it is conducted. On the other hand, it is a measure of the average reduction in uncertainty after the experiment has been conducted. Putting this latter interpretation slightly differently, entropy is the average increase in certainty as a result of conducting the experiment.

In the usual epistemic context, uncertainty results because of a lack of information. We may also define notions of uncertainty in the praxeic context. Even if the decision maker is epistemically certain that a given option is the correct choice, it still may be

equivocal about how well, relative to the other available options, the given one will perform both in terms of achieving the goal and in terms of consuming resources. To illustrate, suppose X can choose among three routes, r_1 , r_2 , and r_3 to take from home to work, and desires to both enjoy the scenery and keep travel time down. On a scale of 0–10, X ranks the scenic beauty of these routes as 4, 5, and 7, with corresponding travel times are 20, 30, and 50 minutes, respectively. Viewing scenic beauty as selectable and travel time as rejectable, the corresponding mass functions are

$$p_S(r_1) = 0.250, \quad p_S(r_2) = 0.312, \quad p_S(r_3) = 0.438,$$

and

$$p_R(r_1) = 0.2, \quad p_R(r_2) = 0.3, \quad p_R(r_3) = 0.5.$$

The variations in the degree of scenic beauty among the options generates a type of equivocation, since each route offers some degree of scenic beauty. Also, the variation in time of travel generates another type of equivocation with respect to consuming resources (time). These equivocations are manifestations of praxeic uncertainty. This type of uncertainty deals with the difficulty the agent has in making a decision. If, for example, all of the selectability mass were ascribed to, say, r_1 , then the praxeic uncertainty with respect to selectability would be negligible. Thus, just as entropy provides a measure of epistemic uncertainty when there is a lack of information, entropy may also provide measures of praxeic uncertainty when the decision maker is equivocal about its choices. Praxeic entropy provides a measure of the “emotional” difficulty the decision maker experiences in making choices.

To appreciate entropy in the satisficing context, we require interpretations of this notion for both selectability and rejectability that are analogous to the usual Shannon interpretation. Let us first consider selectability, and view the expediency of an option as the degree to which it leads to success. Then *inexpediency*, the degree to which an option fails to achieve the goal, is analogous to epistemic uncertainty, or the degree to which an outcome is unlikely to occur. If p_S is a selectability mass function and $p_S(u) \approx 1$, then $-\log_2 p_S(u) \approx 0$, which indicates that, since implementing u is highly selectable, there is little inexpediency associated with doing so. Conversely, suppose $p_S(u') \approx 0$. Then $-\log_2 p_S(u')$ is large, indicating that great inexpediency is associated with implementing that option. $H(p_S)$ then becomes a measure of the average inexpediency associated with the decision problem before taking action. Equivalently, it is a measure of the average reduction in inexpediency after taking action, or to put it more positively, it is the average increase in the degree to which success will be achieved as a result of taking action.

To interpret the entropy of p_R , let us view expense as the degree to which resources are consumed, and consider *inexpense*, the degree to which we conserve resources, as analogous to uncertainty. If $p_R(u) \approx 1$, then $\log_2 p_R(u) \approx 0$, which indicates that u is highly rejectable, thus little inexpense obtains if a highly rejectable option is implemented. On the other hand, if $p_R(u') \approx 0$ and u' is implemented, then $-\log_2 p_R(u')$ is large, indicating that great inexpense obtains if u' is implemented. $H(p_R)$ is a measure of the average inexpense associated with the

decision problem before taking action. Equivalently, it is a measure of the average reduction in inexpense after taking action, or to put it more positively, it is the average increase in the degree to which resources are consumed as result of taking action.

Entropy is maximized by the uniform distribution. Let n be the cardinality of the option space, U (assumed to be finite). If $p^*(u) = 1/n$ for all $u \in U$, then $H(p^*) \geq H(p)$ for all mass functions p over U , and $H(p^*) = \log_2 n$. A near-uniform p_S would generate high average inexpediency, in that all options would work equally (either effectively or ineffectively). A low-entropy p_S would indicate that most of the selectability mass is concentrated on a few options that are highly conducive to success. A near-uniform p_R would generate high average inexpense, in that all of the options cost the same, and none are inexpensive, while a low-entropy p_R indicates that the rejectability mass is concentrated on a few options that consume a disproportionate amount of resource (and, consequently, there exists a subset of options that are inexpensive, in that implementing them will conserve resources). For the drive-to-work example defined above, $H(p_S) = 1.548$, $H(p_R) = 1.486$, and $H(p^*) = 1.585$. Thus, we see that there is considerable praxeic uncertainty with that decision problem. This is reflected in the fact that, with $q = 1$, the satisficing set is $\Sigma_1 = \{r_1, r_2\}$, and there is no obvious way to prefer one route to the other. Although X has no epistemic uncertainty (all routes will get X to work), X does have a non trivial “emotional” decision to narrow the choice to a single option.

Definition 8: If $p_S(u) > 1/n$ (that is, selectability under p_S is greater than selectability under the uniform distribution), then the option is attractive with respect to performance— u is *expedient*.

Definition 9: If $p_R(u) > 1/n$ (that is, rejectability under p_R is greater than rejectability under the uniform distribution), then u is unattractive with respect to resource consumption— u is *expensive*.

The relationship between selectability and rejectability permits the definition of four dispositional modes of the decision maker with respect to each of its options.

Definition 10: If $u \in U$ is both expedient and expensive, then the decision maker will be in a position of desiring to reject, on the basis of cost, an option that is suitable in terms of performance—it will be *ambivalent* with respect to u .

Definition 11: If $u \in U$ is both inexpedient and inexpensive, then the decision maker will be desirous of accepting the option on the basis of cost, but will be reluctant to do so because of poor performance. The decision maker will be *dubious* with respect to u .

Definition 12: If $u \in U$ is expedient and inexpensive, then the decision maker is in the position of desiring, on the basis of cost, to implement an option that would also yield good performance—it will be *gratified* with respect to u .

Definition 13: If $u \in U$ is inexpedient and expensive, then the decision maker will desire to reject, on the basis of cost, an option that also provides poor performance, and will thus be in mode of *relief* with respect to u because it will not be chosen.

These four modes provide a qualitative way for a decision maker to evaluate its choices more definitively. Gratification and

relief are modes of contentment, while dubiety and ambivalence are modes of conflict. With the drive-to-work example, X is dubious with respect to both satisficing decisions r_1 and r_2 .

With multi-agent decision problems, it is in conflictive situations that negotiation is the most difficult. By categorizing the options according to these modes, the decision maker may invoke additional situation-dependent criteria, such as task urgency or resource reserves, to facilitate a compromise. These modes provide additional insight into the process of negotiation, and serve as indicators of the difficulty or ease of making compromise choices.

V. SINGLE-ROUND NEGOTIATIONS

The BOS game described in Section I provides a simple example of single-round negotiation, such as occurs when there are only two players and each player must choose between two alternatives. To see how satisficing methodology might apply to this situation, let us cast BOS as a satisficing game. Although we will retain the traditional story-line, it is easy to adopt this example to a different context. We must first establish each player’s notions of selectability and rejectability. Although there are many ways to frame this game, let us take selectability as the two players being with each other, regardless of where they go. This is the fundamental goal. The resources available to the players are the venues they may attend; rejectability deals with the costs of being at a particular venue. The dichotomy before the players is that the fundamental goal of being together potentially conflicts with the preference for one’s favorite venue. According to the stereotypical roles of the players, H would prefer D if he did not take into consideration S ’s preferences; similarly, S would prefer B . Thus, we may express the myopic rejectabilities for H and S in terms of parameters h and s , respectively, as

$$\begin{aligned} p_{R_H}(D) &= h \\ p_{R_H}(B) &= 1 - h \end{aligned} \quad (9)$$

and

$$\begin{aligned} p_{R_S}(D) &= 1 - s \\ p_{R_S}(B) &= s \end{aligned} \quad (10)$$

where h is H ’s rejectability of D and s is S ’s rejectability of B . The closer h is to zero, the more H is adverse to B with an analogous interpretation for s with respect to S attending D . To be consistent with the stereotypical roles, we may assume that $0 \leq h < 1/2$ and $0 \leq s < 1/2$. As will be seen subsequently, only the ordinal relationship need be specified, that is, either $s < h$ or $h < s$.

Since being together is a joint, rather than an individual, objective, it is difficult to form unilateral assessments of selectability, but it is possible to characterize individually the conditional selectability. To do so requires the specification of the conditional mass functions $p_{S_H|R_S}$ and $p_{S_S|R_H}$, that is, H ’s selectability conditioned on S ’s rejectability and S ’s selectability conditioned on H ’s rejectability. If S were to place her entire unit mass of rejectability on D , H may account for this, if he cares

about S 's feelings, by placing some portion of his conditional selectability mass on B . By so doing, H is manifesting situational altruism. S may construct her conditional selectability in a similar way, yielding

$$\begin{aligned} p_{S_H | R_S}(D | D) &= 1 - \alpha \\ p_{S_H | R_S}(B | D) &= \alpha \\ p_{S_H | R_S}(D | B) &= 1 \\ p_{S_H | R_S}(B | B) &= 0 \end{aligned} \quad (11)$$

and

$$\begin{aligned} p_{S_S | R_H}(D | D) &= 0 \\ p_{S_S | R_H}(B | D) &= 1 \\ p_{S_S | R_H}(D | B) &= \beta \\ p_{S_S | R_H}(B | B) &= 1 - \beta. \end{aligned} \quad (12)$$

The valuations $p_{S_H | R_S}(B | D) = \alpha$ and $p_{S_S | R_H}(D | B) = \beta$ determine the amount of deference one player gives the other. If S were to place all of her rejectability mass on D , then H may defer to S 's strong dislike of D by placing α of his selectability mass, *as conditioned by her preference*, on B . Similarly, S could show a symmetric conditional preference for D if H were to reject B strongly. The parameters α and β serve as a way for each to control the amount of situational altruism he and she are willing to offer. In the interest of simplicity, we shall assume that both players are maximally situationally altruistic and set $\alpha = \beta = 1$. In principle, however, they may be set independently to any value in $[0, 1]$. Notice that, even in this most deferential case, these conditional preferences do not commit one to categorical abdication of his or her own unilateral preferences. H still myopically (that is, without taking S into consideration) prefers D and S still myopically prefers B , and there is no intimation that either participant must "throw" the game in order to accommodate the other.

With these conditional and marginal functions, we may factor the interdependence function as follows:

$$\begin{aligned} p_{S_H S_S R_H R_S}(u_1, u_2; v_1, v_2) \\ = p_{S_H | R_S}(u_1 | v_2) p_{S_S | R_H}(u_2 | v_1) \\ p_{R_H}(v_1) p_{R_S}(v_2) \end{aligned}$$

where we have assumed that H 's selectability conditioned on S 's rejectability is dependent only on S 's rejectability, that S 's selectability conditioned on H 's rejectability is dependent only on H 's rejectability, and that the myopic rejectability values of H and S are independent.

The resulting joint selectability and rejectability functions are

$$\begin{aligned} p_{S_H S_S}(D, D) &= (1 - h)s \\ p_{S_H S_S}(D, B) &= hs \\ p_{S_H S_S}(B, D) &= (1 - h)(1 - s) \\ p_{S_H S_S}(B, B) &= h(1 - s) \end{aligned} \quad (13)$$

and

$$\begin{aligned} p_{R_H R_S}(D, D) &= h(1 - s) \\ p_{R_H R_S}(D, B) &= hs \\ p_{R_H R_S}(B, D) &= (1 - h)(1 - s) \\ p_{R_H R_S}(B, B) &= (1 - h)s. \end{aligned} \quad (14)$$

The marginal selectability and rejectability values for H and S are

$$p_{S_H}(D) = s \quad p_{R_H}(D) = h \quad (15)$$

$$p_{S_H}(B) = 1 - s \quad p_{R_H}(B) = 1 - h \quad (16)$$

and

$$p_{S_S}(D) = 1 - h \quad p_{R_S}(D) = 1 - s \quad (17)$$

$$p_{S_S}(B) = h \quad p_{R_S}(B) = s. \quad (18)$$

Setting the negotiaion index q equal to unity, we obtain the jointly satisficing set as

$$\Sigma_q = \begin{cases} \{(D, B), (B, D), (B, B)\} & \text{for } s < h \\ \{(D, D), (D, B), (B, D)\} & \text{for } s > h \\ \{(D, D), (D, B), (B, D), (B, B)\} & \text{for } s = h \end{cases}$$

the individually satisficing sets are

$$\Sigma_q^H = \begin{cases} \{B\} & \text{for } s < h \\ \{D\} & \text{for } s > h \\ \{B, D\} & \text{for } s = h \end{cases}$$

$$\Sigma_q^S = \begin{cases} \{B\} & \text{for } s < h \\ \{D\} & \text{for } s > h \\ \{B, D\} & \text{for } s = h \end{cases}$$

and the satisficing rectangle is

$$\mathfrak{R}_q = \Sigma_q^H \times \Sigma_q^S = \begin{cases} \{B, B\} & \text{for } s < h \\ \{D, D\} & \text{for } s > h \\ \{\{B, B\}, \{D, D\}\} & \text{for } s = h \end{cases}.$$

Thus, if S 's aversion to D is less than H 's aversion to B , then both players will go to H 's preference, namely, D , and conversely. Notice that these are ordinal, rather than cardinal, comparisons. The satisficing approach fails to give a single answer only in the unlikely situation where both players have exactly equal aversions to the other's preference.

Recall, under classical game theory, that if each player defers to the other, the result is disastrous for both. By contrast, with the satisficing approach, even though both players are maximally conditionally deferential, the satisficing solution is far from disastrous—it results in a very natural cooperative strategy that is socially defensible. Notice, however, that since $s < 1/2$ and $h < 1/2$, by assumption, the dispositional mode of the compromise choice will be dubious for the one who gets to go to his or her favorite venue, and will it be ambivalent for the one who defers. There are no gratifying solutions—the choices are difficult ones for both players. This result provides additional insight for why the BOS game is not easily resolved and why conventional game theory fails to provide a definitive solution.

VI. MULTIROUND NEGOTIATIONS

Many negotiation scenarios are complex, and involve proposals and counter proposals, with each participant modifying its choices and standards for making choices at each round as it seeks for a compromise. Typically, each participant will condition its preferences on the preferences of others. Such conditioning may be based on its own selfish interests, it may be benevolent in the sense of giving deference to others to benefit them at one's own expense, or it may even be malevolent, in the sense of desiring to injure others even if it reduces one's own level of performance. Even in a non-harmonious negotiation scenario, the decision makers may strongly desire to avoid an impasse, especially if the consequences of the group failing to achieve a mutually agreeable decision are high (and thus all players are frustrated) compared to the cost of individuals compromising their individual interests. Thus, there is often an implicit notion of group preference (if only to avoid failure). Such a notion need not be explicitly defined at the outset by the decision makers. Rather, it may emerge as a consequence of their interaction as the conditional preferences propagate through the system.

To illustrate this type of negotiations, we present an example that we name the Three Hermanos (TH). It consists of three agents who act primarily in their own self-interest, but are willing to give some deference to others in order to improve the benefit to the entire group. This scenario involves three brothers, Alberto (A), Juan (J), and Paco (P). They divide their recently deceased father's land into three plots, and each must decide what to grow on his own plot of land for the coming year.

We also assume the following conditions.

- 1) Alberto is the eldest son, receives the best plot of land, and has first choice of which crop to grow. He can grow tomatoes or onions or raise chickens. Juan is the second son, receives the next-best plot of land, and has second choice of which crop to grow. His land will support growing tomatoes, tomatillos, or onions. Paco is the youngest son, has third choice, and the worst plot of land that can grow only beans, tomatillos, or peppers.
- 2) The individual market values for the six possible crops are, in arbitrary units: tomatoes (20), chickens (19), onions (18), beans (17), tomatillos (16), and peppers (15). However, if the brothers cooperate, they can grow products according to three popular recipes with the consequence that they can make and sell these products and thereby increase their income by multiplying the individual market values as follows: enchiladas (chickens, tomatillos, peppers) with a multiplier of $3/2$, burritos (chicken, beans, tomatoes) with a multiplier of $5/4$, and salsa (tomatoes, onions, peppers) with a multiplier of $4/3$.
- 3) The resources required consist of seeds (both for planting the crops and feeding the chickens). In arbitrary units, these costs are: tomatoes (15), chickens (13), onions (14), beans (12), tomatillos (11), and peppers (11). Furthermore, due to the scarcity of these resources, if two brothers decide to plant the same crop, the cost doubles.

This example, although somewhat artificial, nevertheless captures some of the important features of distributed multi-agent

decision making. A centralized approach would be simply to compute the maximum-valued crop and impose that decision upon all of the participants, but that would imply the presence of an external *superplayer* who could dictate the choices to each participant. If such an agent were to exist, the need for negotiation would be obviated. Optimization is instructive, but it is not constructive. It provides a prescription for how individually rational decision makers *should* behave, but does not offer a description of how to achieve the optimal result. Negotiations are required to provide the *process* of making decisions.

A. Conventional Game-Theoretic Approach

To formulate this decision as a game in the tradition of von Neumann and Morgenstern, it would be necessary for each player to specify its utility as a function of possible actions of all players. Since there are three players and each has three options, this means that each player must determine his payoff for each of the possible outcomes. These payoffs would then be juxtaposed in a payoff array, and a solution concept would need to be defined to determine a solution. Although many negotiation protocols exist under the rubric of classical game theory, the requirement that payoffs be defined for all possible outcomes is unwarranted, since they are not specified by the problem statement. Thus, the application of conventional game theory to this problem scenario is problematic.

B. A Satisficing Approach

We begin by identifying the decision spaces for Alberto, Juan, and Paco, respectively, as

$$U_A = \{C, T, O\}$$

$$U_J = \{T, t, O\}$$

$$U_P = \{B, t, P\}$$

where C = chickens, T = tomatoes, O = onions, t = tomatillos, B = beans, and P = peppers.

The next step is to specify operational definitions for selectability and rejectability. We follow the general rule of associating selectability with achieving the fundamental goal of the endeavor, which is to sell the crop, and we take rejectability as being associated with the consumption of resources (seed). Using these operational definitions, we may elicit the following influence relationships.

- 1) Due to the hierarchical nature of the relationships, Alberto's selectability is unconditional.
- 2) Juan's selectability is conditioned on Alberto's selection.
- 3) Paco's selectability is conditioned on both Alberto's and Juan's selections.
- 4) Since growing the same crop greatly increases the cost of resources, the limitation of seed dictates that Juan's rejectability is conditioned on Alberto's selection and that Paco's rejectability is conditioned on Juan's selection. Furthermore, since Alberto and Paco have no crops in common, Alberto's rejectability is conditioned only on Juan's selection.




Fig. 2. Praxeic network for the Three Hermanos.

These influence flows define a praxeic network consisting of six nodes corresponding to the six selves S_A , R_A , S_J , R_J , S_P , and R_P , which correspond to the selectabilities and rejectabilities of Alberto, Juan, and Paco, respectively. The DAG associated with this network is identical with the one given in Fig. 1, where Alberto corresponds to agent 1, Juan to agent 2, and Paco to agent 3. This graph is reproduced with the appropriate relabeling in Fig. 2. The interdependence function corresponding to this network is identical to (8), namely

$$\begin{aligned} p_{S_A S_J S_P R_A R_J R_P}(u_A, u_J, u_P; v_A, v_J, v_P) \\ = p_{S_P | S_A S_J}(u_P | u_A, u_J) p_{R_A | S_J}(v_A | u_J) p_{R_J | S_A}(v_J | u_A) \\ p_{R_P | S_J}(v_P | u_J) p_{S_J | S_A}(u_J | u_A) p_{S_A}(u_A) \end{aligned} \quad (19)$$

where u_A and v_A are elements of U_A , u_J and v_J are elements of U_J , and u_P and v_P are elements of U_P .

The next step in our development is to define the functional values for the six mass functions that appear on the right-hand side of (19). Since Alberto is able to define his selectability unconditionally, he may do so by simply normalizing the market values of the crops available to him, yielding

$$p_{S_A}(C) = \frac{19}{57} \quad p_{S_A}(T) = \frac{20}{57} \quad p_{S_A}(O) = \frac{18}{57}.$$

Next, since Juan's selectability depends upon Alberto's selection, we must define three conditional selectability mass functions $p_{S_J}(\cdot | C)$, $p_{S_J}(\cdot | T)$, and $p_{S_J}(\cdot | O)$. Since Juan is motivated to increase his success, he will desire to choose a crop that will result in one of the recipes. We first consider $p_{S_J}(\cdot | C)$. Given that Alberto selects chickens, Juan may immediately discount salsa, and focus his selectability on tomatillos and onions (to complete the recipe for enchiladas or burritos). Thus, he should ascribe selectability to these choices according to the nominal return on the product he grows multiplied by the appropriate multiplier for the recipe. The selectability mass function corresponding to this logic is

$$\begin{aligned} p_{S_J | S_A}(T | C) = \frac{25}{49} \quad p_{S_J | S_A}(t | C) = \frac{24}{49} \\ p_{S_J | S_A}(O | C) = 0. \end{aligned}$$

To compute Juan's conditional selectability given that Alberto selects tomatoes, we observe that there is only one possible

recipe, hence Juan must ascribe his entire conditional selectability mass to onions, yielding

$$p_{S_J | S_A}(T | T) = 0 \quad p_{S_J | S_A}(t | O) = 0 \quad p_{S_J | S_A}(O | T) = 1.$$

By similar logic, we obtain

$$p_{S_J | S_A}(T | O) = 1 \quad p_{S_J | S_A}(t | O) = 0 \quad p_{S_J | S_A}(O | O) = 0.$$

The conditional selectabilities for Paco, given the choices of Alberto and Juan, are found by similar logic, except that Paco must condition on both Alberto's and Juan's selections. This requires a total of nine conditional selectability mass functions, but they are of simple structure. For example, if Alberto selects tomatoes and Juan selects onions, then Paco should ascribe his entire selectability mass to peppers. The other conditional selectabilities are determined similarly.

The next task is to determine the rejectabilities of the three brothers. We illustrate how this is done by examining Juan's situation; the cases for Paco and Alberto are similar. We need to compute $p_{R_J | S_A}$. Suppose Alberto chooses chickens. Then there can be no conflict, so Juan simply normalizes the seed costs for his three possible crops, yielding

$$\begin{aligned} p_{R_J | S_A}(T | C) = \frac{15}{40} \quad p_{R_J | S_A}(t | C) = \frac{11}{40} \\ p_{R_J | S_A}(O | C) = \frac{14}{40}. \end{aligned}$$

If, however, Alberto were to select tomatoes, then Juan must double the cost of tomato seed for himself, yielding, after normalization,

$$\begin{aligned} p_{R_J | S_A}(T | T) = \frac{30}{55} \quad p_{R_J | S_A}(t | T) = \frac{11}{55} \\ p_{R_J | S_A}(O | T) = \frac{14}{55}. \end{aligned}$$

By a similar calculation, if Alberto were to select onions, then Juan's conditional selectability would become

$$\begin{aligned} p_{R_J | S_A}(T | O) = \frac{15}{54} \quad p_{R_J | S_A}(t | O) = \frac{11}{54} \\ p_{R_J | S_A}(O | O) = \frac{28}{54}. \end{aligned}$$

The conditional rejectability functions for Paco and Alberto can be computed similarly.

These conditional mass functions define the links between the nodes of the praxeic network. Once these links are forged, the joint and individual selectability and rejectability functions can be computed. To do so, however, we must first define initial values the negotiation indices for the group and for the individual. Let q_A , q_J , and q_P denote the negotiation indices for Alberto, Juan, and Paco, respectively, and let q_G denote the negotiation index for the group (for example, we could take $q_G = \min\{q_A, q_J, q_P\}$). Then the joint and individual satisfying sets are

$$\begin{aligned} \Sigma_{q_G} = \{(u_A, u_J, u_P) : p_{S_A S_J S_P}(u_A, u_J, u_P) \\ \geq q_G p_{R_A R_J R_P}(u_A, u_J, u_P)\} \end{aligned}$$

TABLE II
NEGOTIATION OUTCOMES FOR THE THREE HERMANOS PROBLEM

round	q values	Alberto	Juan	Paco	Group
1	(1.0, 1.0, 1.0)	C	T	P	\emptyset
2	(1.0, 1.0, 0.9)	C	T	P	\emptyset
3	(1.0, 0.9, 0.9)	C	T, O	P	\emptyset
4	(0.9, 0.9, 0.9)	C	T, O	P	\emptyset
5	(0.9, 0.9, 0.8)	C	T, O	P	\emptyset
6	(0.9, 0.8, 0.8)	C	T, O, t	P	C, t, P

and

$$\Sigma_{q_A} = \{u_A : p_{S_A}(u_A) \geq q_A p_{R_A}(u_A)\}$$

$$\Sigma_{q_J} = \{u_J : p_{S_J}(u_J) \geq q_J p_{R_J}(u_J)\}$$

$$\Sigma_{q_P} = \{u_P : p_{S_P}(u_P) \geq q_P p_{R_P}(u_P)\}.$$

The *satisficing rectangle* is then

$$\mathfrak{R} = \Sigma_{q_A} \times \Sigma_{q_J} \times \Sigma_{q_P}$$

and the compromise set is $\mathcal{C} = \Sigma_{q_G} \cap \mathfrak{R}$. If $\mathcal{C} = \emptyset$, then there is no set of satisficing options that are satisfying for both the group and for all of the individuals, and some negotiation must take place in order to avoid an impasse.

There are a number of negotiation protocols that could be implemented, with one of the simpler being a round-robin procedure of decrementing the negotiation indices of the three participants. So doing will enlarge the individual satisficing sets, which will in turn enlarge the satisficing rectangle. Eventually, the compromise set will not be empty, and the negotiations can be successfully concluded. A reasonable protocol that is consistent with the hierarchical structure of this society is for Paco, the lowest ranking member, to lower his standards by first decrementing q_P by a small amount, say Δq , resulting in an enlarged satisficing rectangle. If the resulting compromise set is still empty, then Juan would decrement q_J . If the compromise set still remains empty, then Alberto would decrement q_A . If additional rounds or negotiation are required, then Paco would decrement q_P again, and the process would continue until $\mathcal{C} \neq \emptyset$ or until no participant is willing to further lower his q value, resulting in an impasse. If \mathcal{C} contains exactly one set of options, that set is then implemented. If \mathcal{C} contains more than one set of options, then the set that results in the greatest joint selectability for the group would be an appropriate choice.

For the Three Hermanos problem as described, with $\Delta q = 0.1$, six rounds of negotiation were required in order to achieve a compromise. The outcomes of these rounds are summarized in Table II, where the last four columns indicate the individual and group satisficing sets. These results indicate that both Alberto and Paco hold firm to their favorite choice, and it is Juan who eventually gives in to allow the compromise. The resulting jointly and individually satisficing sets are given in Table III. The satisficing rectangle is

$$\mathfrak{R} = \{\{C, T, P\}, \{C, O, P\}, \{C, t, P\}\}$$

and the compromise set is therefore

$$\mathcal{C} = \{C, t, P\}.$$

TABLE III
FINAL JOINTLY AND INDIVIDUALLY SATISFICING VALUES FOR THE THREE HERMANOS PROBLEM

Jointly Satisficing Sets				
Crops	$p_{S_A S_J S_P}$	$p_{R_A R_J R_P}$	Group Attitude	Group Profit
$\{T, O, P\}$	0.351	0.092	Ambivalent	30.7
$\{C, t, P\}$	0.163	0.018	Gratified	40.0
$\{O, T, P\}$	0.316	0.086	Ambivalent	30.7
$\{C, T, B\}$	0.170	0.010	Gratified	30.0

Individually Satisficing Sets				
Brother	Choice	p_S	p_R	q
Alberto	C	0.333	0.203	0.9
Juan	T, O, t	0.486, 0.351, 0.163	0.420, 0.382, 0.198	0.8
Paco	P	0.830	0.462	0.8

Hence the satisficing cash crop is for Alberto to raise chickens, Juan to grow tomatillos, and Paco to grow peppers; the brothers can then combine these ingredients to make enchiladas. Notice that every jointly satisficing option vector is a recipe. Furthermore, it turns out that the compromise option vector is the one that maximizes profits for the group. This result was not stated as an explicit goal of the negotiation; rather, it *emerged* as a group preference as a result of the conditional preferences propagating through the system via the chain rule. This result was not guaranteed. It obtained because the participants were willing to lower their individual standards. If they had not been willing to do so, they would not have achieved the optimal solution. Hence, the group benefits because the individuals are willing to be flexible in their choices, and are not intransigent utility maximizers. A willingness to be moderate at the local (individual) level can turn out to be instrumentally optimal at the global (group) level. The success of this negotiation is evidenced by the fact that the compromise solution is gratifying for the group.

The individual dispositional modes associated with the compromise decision are that Alberto is gratified, Juan is dubious, and Paco is ambivalent. These modes help interpret the negotiation process. During the negotiation, neither Alberto nor Paco were willing to budge from their initial choice. Juan, however, was the one who made possible the compromise. The fact that Juan's choice is dubious (both selectability and rejectability are very low) helps to explain this situation.

VII. DISCUSSION

This paper presents a negotiation theory that is based on a formalized game-theoretic structure that is as mathematically rigorous as is conventional game theory as developed by von Neumann and Morgenstern. Von Neumann–Morgenstern game theory is based on the hypothesis of individual rationality, and is therefore of limited value for situations where cooperative negotiations are essential. Since satisficing game theory is not founded on individual rationality, it is able to accommodate both non-cooperative and cooperative negotiation scenarios. This paper has focused primarily on the cooperative situation, since that is the scenario where the power of this approach is perhaps most obvious. However, the satisficing theory is applicable to general negotiation scenarios. Some of the features that apply to the more general case include the following.

- 1) Since negotiation protocols are distributed, it is not required or assumed that all participants subscribe to the satisficing point of view. Even if a satisficer is negotiating with a non-satisficer, it may proceed according to a protocol based on satisficing game theory. By so doing, the agent is able to identify all options for itself that are good enough as defined by its criteria, and it is able to control the degree to which it is willing to lower its standards as it attempts to achieve a compromise. Satisficing is more flexible than optimization. It provides some friction to the slippery slope of compromise.
- 2) Whereas optimization is strictly an individual concept, satisficing can be a social, as well as an individual, concept. For any group of decision makers, if the group and each of its members are willing to compromise sufficiently (either out of deference to others or simply because the penalty for failing to reach a compromise is catastrophic), there will exist a joint option that is good enough for the group as a whole and good enough for each member of the group according to their individual standards. This does not mean, of course, that the decision makers are obligated to accept this compromise option. It means only that it exists.
- 3) The interdependence function is able to accommodate self-interest as well as community interest. Therefore, a self-interested player is able to encode exactly the same information using satisficing theory as can be done via von Neumann–Morgenstern utilities. In fact, satisficing theory is even more general than individual rationality in that it permits situational altruism. We do not assert that, under the theoretical framework of conventional game theory, it is impossible to formulate theoretical models of social behavior that go beyond individual interests and accommodate situationally altruistic tendencies while at the same time preserving individual preferences. However, the extant literature does not provide such a theory.
- 4) Virtually all negotiation protocols provide rationale for an agent to modify its position in order to seek a mutually acceptable solution. Usually, such procedures require the agent to change from its most preferred outcome to an outcome that is less preferred. Since conventional game theory requires the decision maker to do the best for itself, using it as a protocol for negotiation requires some mechanism for the agent to revise its utilities. Such mechanisms are not part of the basic theory, and there is no systematic way to introduce them without making additional assumptions that are not part of the game-theoretic structure. Satisficing, on the other hand, provides a systematic and convenient mechanism for the agents to modify their standards; namely, they may iteratively adjust q , their negotiation indices. By so doing, they gradually widen their consideration of alternatives in a controlled way. They may set explicit limits as to how far they are willing to go in order to accommodate others, and they may break off negotiations if they would be required to sacrifice more performance than they can afford. Notice that this form of compromise does not require the agent to modify its preferences as expressed by its utilities. Rather, it only requires it to modify the negotiation index.
- 5) Regardless of the notions of rationality, the negotiation protocols, or any other aspects of negotiation, the success of any negotiating agent is limited by the accuracy of its model of the environment. Consequently, an important aspect of any negotiation protocol is the ability to learn, and satisficing theory provides a particularly convenient way to accommodate this requirement. Recall that the interdependence function is formed as the product of conditional selectability and rejectability functions, each of which corresponds to a hypothetical situation involving the preferences of other agents. By actualizing such hypothetical situations, a satisficing negotiator can learn the preferences of the other players and thereby dynamically adapt its interdependence function to the actual situation. It is beyond the scope of this paper, however, to develop such learning procedures in detail.
- 6) Two types of complexity arise with the satisficing approach: a) modeling complexity and b) computational complexity. Extending the sphere of interest beyond the self increases the complexity of a multiagent system model, since it must account for sophisticated social relationships such as compromise, negotiation, and altruism. As noted by Palmer, “Complexity is no argument against a theoretical approach if the complexity arises not out of the theory itself but out of the material which any theory ought to handle” [27]. If one is to account for social relationships that exist between members of a multi-agent system, one must pay the price. Computational complexity arises because of the calculation of the marginals of the interdependence function. This complexity can be mitigated somewhat by efficient organization of the computations, using, for example, Pearl’s Belief Propagation Algorithm [25] or the factor graph approach described by [26]. Even so, it is well known that even these approaches are *NP* hard, and the computational burden for a tightly interconnected, high-dimensional multi-agent system may become intractable. Fortunately, however, as is the case with many useful Bayesian networks, many interesting multi-agent systems will be rather sparsely connected.
- 7) A possible weakness of the satisficing approach is that, by eschewing optimality as the ideal, the participants may settle on a “good enough” solution that is of dubious quality, that is, one for which neither the benefits nor the costs are very high. The fact of the matter is, however, that neither optimization nor satisficing can guarantee that the chosen solution is very good. Making the best of a bad situation is not very comforting, but at least with the satisficing approach, the players are able to evaluate the quality according to attitudinal modes.
- 8) Satisficing negotiation provides an explanation for some forms of human negotiation. In some negotiations (for example, the TH problem), parties often repeat a position without an apparent change of state, until at some point there is an abrupt change in feasible options. The

procedure presented here provides a model for such behavior: even when from one iteration to the next there might be no change in the compromise sets, each agent is lowering its negotiation index and, if it possesses a learning ability, is modifying its models of the other agents. Other behaviors, such as recalcitrance or openness, can be accommodated by more sophisticated dynamics of how the negotiation indices are changed.

VIII. CONCLUSION

Multi-agent satisficing theory provides a means of describing solutions which are individually and jointly satisficing from the perspective of an individual agent in a community of agents. This paper extends previous work on satisficing decision theory by a) distinguishing between categorical and situational altruism, b) providing a discussion of participant attitude modes, c) providing an explicit protocol for negotiation under the satisficing regime, and d) providing examples that demonstrate negotiation under the satisficing paradigm.

Satisficing game theory provides a new tool for the analysis and design of multi-agent systems. It is particularly applicable to negotiatory situations since, by substituting a mathematically precise notion of being good enough for the notion of optimality, it provides decision makers with flexibility to adjust their choices as they interact with each other. The theory is equally applicable to cooperative and noncooperative scenarios.

When cooperation is essential in a multi-agent system, it is important to design the system according to principles that explicitly accommodate cooperation. However, under conventional individual rationality-based approaches such as von Neumann–Morgenstern game theory, it is difficult to characterize cooperation, especially if it requires deferring preferences at one's own expense in order to benefit others. But under the notion of satisficing rationality, giving deference is easy to characterize and to specify via conditional preference relationships.

The appeal of optimization is a strongly entrenched attitude that dominates current decision-making practice. There is great comfort in following traditional paths, especially when those paths are founded on such a rich and enduring tradition as individual rationality affords. But when synthesizing an artificial negotiatory system, the designer has the opportunity to impose upon the agents a more socially accommodating paradigm. The satisficing game theory presented in this paper provides a sociological decision-making mechanism that seamlessly accounts for group and individual interests, and provides a rich framework for negotiation to occur between agents who share common interests and who are willing to give deference to each other. Rather than depending upon the noncooperative equilibria defined (even if only approximately) by individual benefit, this alternative may lead to the more socially realistic and valuable equilibria of shared interests and acceptable compromises.

REFERENCES

- [1] J. von Neumann and O. Morgenstern, *The Theory of Games and Economic Behavior*. Princeton, NJ: Princeton Univ. Press, 1944. 2nd ed., 1947.
- [2] T. W. Sandholm, "Distributed rational decision making," in *Multiagent Systems*, G. Weiss, Ed. Cambridge, MA: MIT Press, 1999, ch. 5, pp. 201–258.

- [3] S. J. Brams, *Negotiation Games: Applying Game Theory to Bargaining and Arbitration*. New York: Routledge, 1990.
- [4] A. Bergson, "A reformulation of certain aspects of welfare economics," *Quart. J. Econ.*, vol. 52, pp. 310–334, 1938.
- [5] P. A. Samuelson, *Foundations of Economic Analysis*. Cambridge, MA: Harvard University Press, 1948.
- [6] M. Taylor, *The Possibility of Cooperation*. Cambridge: Cambridge Univ. Press, 1987.
- [7] K. J. Arrow, *Social Choice and Individual Values*. New York: John Wiley, 1951. 2nd ed., 1963.
- [8] R. D. Luce and H. Raiffa, *Games and Decisions*. New York: John Wiley, 1957.
- [9] H. A. Simon, "A behavioral model of rational choice," *Quart. J. Econ.*, vol. 59, pp. 99–118, 1955.
- [10] —, "Rational choice and the structure of the environment," *Psych. Rev.*, vol. 63, no. 2, pp. 129–138, 1956.
- [11] G. Hardin, "The tragedy of the commons," *Science*, vol. 162, pp. 1243–1248, 1968.
- [12] R. Nozick, *The Nature of Rationality*. Princeton, NJ: Princeton Univ. Press, 1993.
- [13] K. J. Arrow, "Economic theory and the hypothesis of rationality," in *Utility and Probability*, J. Eatwell, M. Milgate, and P. Newman, Eds. New York: W. W. Norton and Company, 1990, pp. 25–37.
- [14] W. C. Stirling and D. R. Morrell, "Convex bayes decision theory," *IEEE Trans. Syst., Man, Cybernet.*, vol. 21, no. 1, pp. 173–183, Jan./Feb. 1991.
- [15] W. C. Stirling, "Multi-agent coordinated decision-making using epistemic utility theory," in *Artificial Social Systems*, E. Castelfranchi and E. Werner, Eds. Berlin, Germany: Springer-Verlag, 1994, pp. 164–183.
- [16] M. A. Goodrich, "A theory of satisficing control," Ph.D. dissertation, Brigham Young University, Provo, UT, 1996.
- [17] M. A. Goodrich, W. C. Stirling, and R. L. Frost, "A satisficing approach to intelligent control of nonlinear systems," in *Proceedings of the IEEE Symposium on Intelligent Control*, Dearborn, MI, Sep. 1996, pp. 248–252.
- [18] —, "A theory of satisficing decision and control," *IEEE Trans. Syst., Man, Cybernet.*, vol. 28, no. 6, pp. 763–779, 1998.
- [19] W. C. Stirling and M. A. Goodrich, "Satisficing games," *Inf. Sci.*, vol. 114, pp. 255–280, Mar. 1999.
- [20] M. A. Goodrich, W. C. Stirling, and E. R. Boer, "Satisficing revisited," *Minds and Machines*, vol. 10, pp. 79–109, Feb. 2000.
- [21] W. C. Stirling, M. A. Goodrich, and D. J. Packard, "Satisficing equilibria: A non-classical approach to games and decisions," *Auton. Agents Multi-Agent Syst.*, vol. 5, pp. 305–328, 2002.
- [22] W. C. Stirling, "Games machines play," *Minds Machines*, vol. 12, pp. 327–352, 2002.
- [23] —, "Beyond optimization: Overcoming the limitations of individual rationality," in *Proc. Game Theoretic and Decision Theoretic Agents (GTDT'02): Papers from the AAI Workshop Technical Report WS-02-06*, Edmonton, AB, Canada, 2002, pp. 99–108.
- [24] —, *Satisficing Games and Decision Making: With Applications to Engineering and Computer Science*. Cambridge, U.K.: Cambridge University Press, 2003.
- [25] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA: Morgan Kaufmann, 1988.
- [26] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 498–519, 2001.
- [27] F. R. Palmer, *Grammar*. Harmondsworth, Middlesex, U.K.: Harmondsworth, Penguin, 1971.



James K. Archibald (M'00) received the B.S. degree in mathematics from Brigham Young University, Provo, UT, in 1981 and the M.S. and Ph.D. degrees in computer science from the University of Washington in 1983 and 1987, respectively.

Since 1987, he has been with the Electrical and Computer Engineering Department at Brigham Young University. His research interests include robotics and multiagent systems.

Dr. Archibald is a Member of the ACM and Phi Kappa Phi.



Jared C. Hill received the B.S. degree in electrical engineering from Brigham Young University, Provo, UT, in 2005. He is currently pursuing the M.S. degree at Brigham Young University, where he is actively pursuing research in the area of multiagent decision theory.

His educational interests include software engineering, signal processing, and decision theory.



F. Ryan Johnson received the B.S. degree in computer engineering from Brigham Young University, Provo, UT, in 2004. He is currently a graduate student in the ECE Department at Carnegie Mellon University, Pittsburgh, PA, with interests in parallel embedded architectures.

His interests include computer architecture, embedded systems, and algorithms.



Wynn C. Stirling received the Honors B.A. degree (*magna cum laude*) in mathematics and the M.S. degree in electrical engineering from the University of Utah, Salt Lake City, in 1969 and 1971, respectively. He received the Ph.D. in electrical engineering from Stanford University, Stanford CA, in 1983.

From 1972 to 1975, he was with Rockwell International Corporation, Anaheim, CA, and from 1975 to 1984, he was employed by ESL, Inc., Sunnyvale, CA, where he was responsible for the development of multivehicle trajectory reconstruction capabilities.

He joined the faculty of the Department of Electrical and Computer Engineering at Brigham Young University, Provo, UT, in 1984 and holds the rank of Professor. His current research interests include multi-agent decision theory, estimation theory, information theory, and stochastic processes. He is the author or coauthor of over 70 publications. He is coauthor of a Prentice-Hall graduate-level text entitled *Mathematical Methods and Algorithms for Signal Processing* (2000), and is the author of a Cambridge University Press monograph entitled *Satisficing Games and Decision-Making: With Applications to Engineering and Computer Science* (2003). He has served on the program committees for conferences on imprecise probability theory and multi-agent decision theory.

Dr. Stirling is a member of Phi Beta Kappa and Tau Beta Pi.