

Final Report

Investigator: **Arthur Robert Calderbank**
Project Title: **Coding Theory Information Theory and Radar**
Location of Project: **Princeton University**

Source of Support: **Air Force Office of Scientific Research**
Total Award period Covered: **07/15/04 to 07/14/05**
Total Award Amount: **\$26,200**

1. Participants

- **Robert Calderbank** is an expert on methods of limiting the end to end complexity of signal processing for multiple antenna systems so that it is close to that of single antenna systems. In collaboration with Dr. Howard and Dr. Moran, he has provided rigorous methods for the design of libraries of waveforms, or more generally libraries of radar modalities, for detection, identification and tracking application. When different waveforms can be transmitted independently and coherently from different array elements, waveform selection is based on a generalization of the standard radar ambiguity function that describes the response of the radar to a point target at a fixed range and Doppler.
- **Stephen Howard (DSTO, Australia) and Bill Moran (University of Melbourne)** have been involved in the area of electronic surveillance and radar for the past ten years. Within the Defence Science and Technology Organisation Dr. Howard has led the research effort into the development of algorithms in all areas of electronic surveillance, including radar pulse train deinterleaving, precision radar parameter estimation and tracking, estimation of radar intra-pulse modulation and advanced geolocation techniques. Since 2003, he has led the DSTO long range research program in radar resource management and waveform design. In collaboration with Dr. Moran, he has provided rigorous methods for the design of libraries of waveforms, or more generally libraries of radar modalities, for detection, identification and tracking application.

2. Activities and Findings:

- The major research focus was the relationship between new ideas in classical and quantum error correcting codes and the development of an information theory for radar.
- We found that the m -dimensional discrete Heisenberg-Weyl group provides a unifying framework for a number of important sequences significant in the construction of phase coded radar waveforms, in communications as spreading sequences, and in the theory of error correcting codes [1, 2]. Among the sequences which can be associated with the Heisenberg-Weyl group are the first and second order Reed-Muller codes, the Golay or Welti sequences, and the

20050616 163

Kerdock and Preparata codes [7], which are non-linear binary error correcting codes containing more codewords for a given minimum distance than any linear code. The Kerdock codes are associated with decomposition of the Heisenberg-Weyl group into disjoint maximally commutative subgroups [1, 2, 6, and 7]. It is a surprising fact that a certain general class of Golay sequences exist within the Kerdock codes. This had previously been noted by Davis and Jedwab [8].

Golay sequences [3–5, 8] are pairs of sequences of unimodular complex numbers with the property that the sum of their individual auto-correlation functions forms delta spike or thumb tack. These sequences have found application in the construction of radar waveforms and in modulation schemes for communications.

- This project led to the creation of a new cross-functional team (Robert Calderbank, Stephen Howard (DSTO, Australia, and Bill Moran (Melbourne) spanning the application domain, classical harmonic analysis, algebra and combinatorics. The capabilities developed by this team as part of this project have been integrated into a larger multi-institutional response to the task specific DARPA BAA 05-19.
- The results of this project were shared with the mainstream signal processing community at ICASSP'05 (Philadelphia) and with the defense signal processing community at DASP'05 (Homestead, Utah).

References

- [1] S. D. Howard, A. R. Calderbank, W. Moran, H. A. Schmitt and C. O. Savage. Relationships Between Radar Ambiguity and Coding Theory, Proc. IEEE International Conference on Acoustics, Speech and Signal Processing, 2005, March 2005.
- [2] S. D. Howard, A. R. Calderbank and W. Moran. The Discrete Finite Heisenberg Group in Radar and Communications, to appear in EURASIP special issue, 2005.
- [3] M. J. E. Golay, Complementary Series IRE Transactions on Information Theory, 7(12):82–87, April 1961.
- [4] G. R. Welter Quaternary codes for pulsed radar IRE Trans. Inf. Theory, 6:400-408, 1960.
- [5] S. Z. Budisin, B. M. Popovic and I. M. Indjin, Designing Radar Signals Using Complementary Sequences. IEEE Transactions on Aerospace and Electronic Systems 21(2):170–9, March 1985.
- [6] AR Hammons, PV Kumar, AR Calderbank, NJA Sloane, and P Sole, The Z₄-linearity of Kerdock, Preparata, Goethals, and related codes. IEEE Transactions on Information Theory, 40(2):301–19, March 1994.
- [7] A. R. Calderbank, P. J. Cameron, W. M. Kantor and J. J. Seidel, Z₄-Kerdock Codes, Orthogonal Spreads, and Extremal Euclidian Line-Sets Proc. London Math. Soc., 3(75):436–480, 1997.

[8] J. A. Davis and J. Jedwab Peak-to-Mean Power Control in OFDM, Golay Complementary Sequences, and Reed-Muller Codes IEEE Trans. Info. Theory, 45(7):2397-2417, November 1999.

3. Publications and Products

[1] S. D. Howard, A. R. Calderbank, W. Moran, H. A. Schmitt and C. O. Savage. Relationships Between Radar Ambiguity and Coding Theory, Proc. IEEE International Conference on Acoustics, Speech and Signal Processing, 2005, March 2005.

[2] S. D. Howard, A. R. Calderbank and W. Moran. The Discrete Finite Heisenberg Group in Radar and Communications, to appear in EURASIP special issue, 2005.

4. Application Value

This project is the starting point for the development of algorithms to significantly expand existing radar functionality. The appropriate advanced hardware functionality exists in the NRL Chesapeake Bay Detachment radar, and at other sites in the US and in Australia. This facility is capable of high levels of temporal and spatial diversity of waveforms and polarization, and the mode of operation can be rapidly scheduled. This project supports development of techniques to fully exploit the diversity and flexibility provided by this advanced functionality through adaptive multidimensional waveform and polarization scheduling based on environment modeling and tracking, as well as multi-dimensional adaptive processing of the returns.

A New Theory of Cooperative Multiagent Systems: Expansions and Applications

A White Paper Submitted to

The Air Force Office of Scientific Research

Wynn C. Stirling*, Richard L. Frost, and James K. Archibald

Electrical and Computer Engineering Department

459CB Brigham Young University, Provo, Utah 84602

{wynn, rickf, jka}@ee.byu.edu

Harold L. Miller, Jr.

Psychology Department

1124 SWKT Brigham Young University, Provo, Utah 84602

harold.miller.byu.edu

May 27, 2005

*Corresponding author, 801.422.7669.

1 Background and Motivation

Decision making in military environments often involves complex relationships between the decision making agents who must cooperate in order to achieve some objective. Typically, decision makers will be concerned both with group performance (e.g., accomplishing the mission) as well as individual performance (e.g., survival, complying with assigned duties). In many realistic scenarios, decision makers will be distributed in space and time and must operate without a central authority. They may have multiple conflicting priorities and will need to proceed without complete information. For such scenarios, an "optimal" solution may not be well defined, since optimal choices by the individuals do not guarantee optimal behavior for the group and vice versa. Even if an "optimal" solution is feasible, it may be based on assumptions that are difficult to justify or on parameter values that are difficult to verify. Nevertheless, the individuals must be capable of achieving an acceptable level of coherent and coordinated behavior even in the face of adverse circumstances.

Satisficing game theory (SGT), which is summarized by Stirling, Frost, and Miller [28] (which accompanies this White Paper as a supplemental document) as well as more completely discussed by Stirling [27], provides a new theoretical foundation for the design and synthesis of artificial multiagent systems and for the understanding of human decision-making processes. This theory stands in contrast to classical decision methodologies that are based on optimization (i.e., utility maximization). It is a mathematically rigorous theory which, although it employs the descriptive term *satisficing*, is very different from the theory of bounded rationality developed by Simon [21, 22], who employs the term to describe decision-making that achieves a heuristically defined aspiration level. Satisficing, in that sense, is an approximation: optimization remains the ideal, but the cost of achieving it is demonstrably prohibitive. Consequently, satisficing à la Simon is philosophically no different from strict optimization.

Satisficing as defined herein is a mathematically rigorous alternative to optimization-based methodologies. Rather than attempting to identify (even approximately) a single "best" solution, it identifies a set of solutions that are "good enough" in a well-defined mathematical sense. The resulting group and individual solutions conform to two fundamental social axioms: (a) *endogeny*, which stipulates that the preferences of a group must emerge through coordination among its individual members; and (b) *social coherence*, which requires that no subset of individuals should be categorically required to forfeit its own welfare in order to benefit the group. These axioms constitute minimal conditions for the reconciliation of group and individual interests. By contrast, optimization-based solutions do not comply with either axiom. This concept of satisficing is explicitly designed to account for the sophisticated social relationships such as coordination, compromise, negotiation, and altruism that may exist in a complex society.

SGT derives its ability to account simultaneously for both group and individual interests from the structure of the utility functions that it employs. These utilities, called *social utilities*, are constructed according to the syntax (but not the usual semantics) of probability theory. Consequently, they permit the modeling of multiagent and multi-attribute behavior in a way that is mathematically analogous to the way probability theory is used to model multivariate random phenomena. In particular, notions of conditioning and independence can be adapted from the traditional epistemological context (i.e., the classification of propositions in terms of belief and knowledge) to a *praxeological* context (the classification of actions in terms of their efficiency and effectiveness).

SGT is relatively new. It has not yet had wide exposure within scholarly communities that

focus on decision theory and application, nor has it yet been subjected to rigorous testing that would show it a viable means of decision making. However, preliminary (and largely unfunded) results indicate that it has promise. It represents a true alternative to both classical decision theory based on expected utility maximization and more recent heuristic approaches such as rule-based systems, both of which have well-known limitations [13, 17, 18, 21, 25, 26, 30].

This white paper describes four potential research agendas that, if successful, will establish SGT as a maturing and useful approach to decision making, especially in group contexts where multiple decision makers must function autonomously. The topics associated with the proposed research agendas are:

- the further development of multi-attribute satisficing decision theory;
- satisficing decision making under conditions of uncertainty and ignorance;
- an information-theoretic analysis of satisficing game theory; and
- psychological evaluations of satisficing decision theory as a model of human decision making.

2 Multi-Attribute Decision Making

Many decision problems involve the consideration of possibly conflicting criteria and, in order to form a decision, it is necessary to construct a tradeoff between the various criteria that results in an acceptable compromise. The multi-attribute decision problem is isomorphic to the multiagent decision problem where each attribute corresponds to a different agent under the constraints that (a) the agents have a common option set and (b) they must all make the same choice. Viewed this way, the multi-attribute decision problem is equivalent in structure to a social choice problem. When discussing multi-attribute decision issues, therefore, the terms “attribute” and “individual” or “agent” may be used interchangeably.

Many algorithms have been devised to address the multi-attribute decision problem, including pair-wise comparisons, ranking methods, rating methods, weighted sum approaches, strength of preference methods, and constraint satisfaction methods. These techniques share two common features. First, they are based on some notion of optimization; that is, they each attempt to identify a “best” solution according to some criterion. Second, they suffer the consequences of Arrow’s impossibility theorem [1], namely, that it is impossible to formulate a ranking of group preferences that is consistent with individual preference rankings. Thus, Arrow’s theorem guarantees that some of the desirable properties of a well-structured decision theory will be violated for each of these methods. For example, pair-wise comparisons are easily shown to violate transitivity and ranking methods commonly violate the irrelevant alternatives condition. Furthermore, although much attention has been devoted to the development of methodologies for selecting optimal weights with which to aggregate disparate criteria, the fundamental arbitrariness of this endeavor can lead to paradoxes and inconsistencies that are difficult to reconcile. As a result, even multi-attribute decision methodologies that purport to be prescriptive (since they are founded on some notion of optimization) are often not descriptive of actual human decision making.

In contrast to optimization based methods, satisficing multiagent decision theory [27] is not founded on optimization. Instead, it is founded on a mathematically formalized concept of being

good enough, and does not require the rank-ordering of preferences. Consequently, Arrow's theorem does not apply and, by means of the theory described in [28], it is indeed possible to reconcile group and individual preferences.

The root of the problem with conventional approaches to multi-attribute decision making is the assumption that individual (single attribute) preferences can be formed independently of the preferences for other individuals (attributes). However, individual preferences may not be independent; that is, the value of a given attribute may depend on the value of another attribute (e.g., when choosing a place to live, the price one is willing to pay for a dwelling may depend on its location relative to one's workplace). Fortunately, satisficing decision theory is explicitly designed to accommodate such conditional relationships. Furthermore, it is also designed to accommodate compromises and negotiations.

2.1 A Satisficing Approach

We may formulate a social choice problem in terms of satisficing decision theory as follows. Let U denote the set of options, and let A denote a collection of n attributes. For example, when choosing a place to live, U might be a list of possible apartments, and A might be a list of attributes such as distance from work, rent, size, age, etc. Associate with the i th individual in the group a selecting process S_i that considers the options with respect to achieving that individual's goal, and a rejecting process R_i that considers the options with respect to resource consumption. These considerations will result in the assignment of conditional and marginal selectability and rejectability functions which will be used to define the interdependence function $p_{S_1 \dots S_n R_1 \dots R_n}$. Once the interdependence function is formed, the joint selectability function $p_{S_1 \dots S_n}$ and a joint rejectability function $p_{R_1 \dots R_n}$ may be formed. With these functions in place, a satisficing social welfare function may be defined as

$$W_q(u) = p_{S_1 \dots S_n}(u, \dots, u) - qp_{R_1 \dots R_n}(u, \dots, u).$$

The distinction between a satisficing social welfare function and classical social welfare functions is that the former is an endogenous aggregation of each individual interests, while the latter are exogenous aggregations, that is, the structure is imposed by external sources. The advantage of an endogenous aggregation is that it guarantees social coherence in the sense that no individual is required to forfeit its individual welfare *in all situations* in order to benefit the group. This condition seems to be a minimal constraint on a collection of individuals in order for them even to be theoretically able to reach a acceptable compromise.

Once W_q has been computed, a solution can be obtained in several ways. Perhaps the simplest is to choose the option that maximizes satisficing social welfare; namely,

$$u^* = \arg \max_{u \in U} W_q(u).$$

A more sophisticated approach would be to define the *social choice set*

$$\Gamma_q = \{u \in U : W_q(u) \geq 0\}.$$

If the social choice set were empty, there would be no good enough choices from the perspective of the group. This would mean that the interests of the individuals are so disparate that there are no actions that provide a net advantage to the group — in other words, the group would be dysfunctional.

Also, the marginal (individual) selectability and rejectability functions p_{s_i} and p_{r_i} may also be extracted from the interdependence function, and the individual satisficing sets may be computed:

$$\Sigma_q^i = \{u \in U: p_{s_i}(u) \geq qp_{r_i}(u)\}.$$

The *consensus set* is the intersection of all individually satisficing sets

$$\mathcal{C}_q = \bigcap_{i=1}^n \Sigma_q^i.$$

If the consensus set is empty, then there are no mutually acceptable alternatives and no outcome can be pleasing to all individuals (i.e., the attributes are so disparate that there is no way to reconcile them). However, if the individuals are willing to reduce q , then eventually both the social choice set Γ_q and the consensus set $\mathcal{C}_q$ will be non-empty.

The *acclimation set*

$$\mathcal{A}_q = \mathcal{C}_q \cap \Gamma_q.$$

comprises all options that are both good for the group as well as the individuals that compose it. If this set is empty, then no choices that serve the interests of the individuals will also serve the interests of the group.

2.2 Key Research Issues

The brief formulation provided above serves as a point of departure for the further development of a comprehensive theory of multi-attribute decision making. In addition to expanding the theory it must be tested, first by simulations and then by psychological experiments in order to evaluate its plausibility and viability as a descriptive model of human multi-attribute decision making. Thus, the research agenda consists of:

- a detailed theoretical analysis of satisficing multi-attribute decision making, using the above approach as a point of departure.
- Computer simulations of multi-attribute decision making using real-world scenarios;
- extension to the combined multiagent multi-attribute case, where each agent is concerned with multiple attributes :(in this situation, the goal is for each agent satisfy its own multi-attribute decision problem in a way that also complies with the interests of the group);
- psychological studies of the applicability of the satisficing model to human multi-attribute decision making (See Section 5).

3 Decision Making Under Uncertainty and Ignorance

Decisions must often be made in an environment of less than complete knowledge. This lack is manifest in two ways: uncertainty and ignorance. Both of these concepts admit degrees of severity, but many statistical theorists recognize a distinction. Common measures of uncertainty are the entropy and the variance of a distribution. These measures, however, are typically taken

with respect to a single distribution which is assumed to be a complete and accurate description of the randomness of the phenomenon. If a single probability distribution is an adequate description of randomness, then the distribution is said to be *precise*.

Ignorance, on the other hand, is manifest by *imprecision*, which arises if there is not a unique probabilistic description of randomness. Imprecision is manifest by a set of distributions, with the size of the set representing the degree of imprecision. Sources of imprecision include (a) vagueness, (b) insufficient resources to gather and process information in order to define precise models, and (c) the impossibility of reaching a consensus for the description of random behavior.

To account realistically for incomplete knowledge, it is necessary to consider both uncertainty and ignorance. To do so, we first consider the precise case and then extend the result to imprecision.

3.1 Uncertainty

The classical approach to decision making under uncertainty is to compute expected utility, where the expectation is taken with respect to a probability distribution of the random components of behavior. In the satisficing context, a naive approach would be to compute the expected interdependence function, from which expected selectability and rejectability functions can be derived. Such an approach, however, would presuppose that the epistemically random components of the model would affect the praxeic components, but not vice versa. Such a constraint would be quite restrictive. For example, suppose the random component is due to the unpredictable activity of an opponent in a hostile engagement. It is possible that if an agent were to take friendly actions, these would modify the actions taken by an otherwise hostile agent. Thus, it would be wise to permit two-way coupling between praxeic and epistemic components. Fortunately, satisficing game theory can be extended to accommodate this more complex model.

Perhaps the most profound attribute of social utilities is that, since they conform to the syntax of probability theory, there need be no *mathematical* distinction between praxeic and epistemic attributes of a system. Furthermore, graph theory provides a language with which to model behavior in a way that emphasizes the local interactions between entities in a compact and efficient way. To construct such a model, let Θ be a set of possible m -dimensional epistemic states of nature that are known only probabilistically, and let S and R denote the selectable and rejectable processes of an n -agent system. We proceed by constructing a graph with $2n + m$ vertices, with m of them corresponding to the epistemic states and n corresponding each to the selecting and rejecting processes. The edges of such a graph correspond to both praxeic and epistemic influence flows.

Consider a decision problem involving a group of two agents X_1 and X_2 , and let $\{u_1, \dots, u_n\}$ denote a finite set of possible actions for the group. The selecting processes are S_1 and S_2 , which view the possible actions in terms of effectiveness in achieving a goal, and the rejecting processes are R_1 and R_2 , which view the possible actions in terms of conserving resources. These four nodes constitute the praxeic component of the decision problem. In addition, suppose there are also components that model the epistemic uncertainty, or randomness, of the group. To illustrate, consider the two-agent directed acyclic graph (DAG) described in Figure 1, consisting of four praxeic nodes S_1 , S_2 , R_1 , and R_2 , and three epistemic nodes θ_1 , θ_2 , and θ_3 . Notice that the praxeic state S_2 influences the epistemic state θ_2 , which in turn affects the praxeic state R_1 , and the epistemic state θ_1 affects the praxeic state R_2 , which affects the epistemic state θ_3 and the praxeic state R_1 . Also, the epistemic state θ_3 affects the praxeic states R_1 and S_1 . Thus, the mathematical structure (using the syntax of probability theory) permits both the epistemic and praxeic components

to be characterized by a joint mass function that accounts for interactions between and within the components.




Figure 1: An example of an epistemi-praxeic DAG.

The epistemi-praxeic interdependence function for the network displayed in Figure 1 is a mass function of the form

$$\begin{aligned}
 p_{S_1 S_2 R_1 R_2 \theta_1 \theta_2 \theta_3}(u_1, u_2, v_1, v_2, \vartheta_1, \vartheta_2, \vartheta_3) = \\
 p_{S_1 | \theta_3}(u_1 | \vartheta_3) p_{S_2}(u_2) p_{R_1 | R_2 \theta_2 \theta_3}(v_1 | v_2, \vartheta_2, \vartheta_3) \\
 p_{R_2 | \theta_1}(v_2 | \vartheta_1) p_{\theta_1}(\vartheta_1) p_{\theta_2 | S_2}(\vartheta_2 | u_2) p_{\theta_3 | R_2}(\vartheta_3 | v_2)
 \end{aligned}$$

This structure seamlessly combines the epistemic (statistically random) components of a multi-agent system with its praxeic (action-taking) components. By applying standard algorithms such as Pearl's Belief Propagation Algorithm [20] it will be possible to compute the marginal expected selectability and rejectability functions, where the expectation is taken over the epistemic variables θ .

3.2 Imprecision

For many applications it will be impossible to specify unique influence relationships between the selecting and rejecting processes. For example, it may be only possible to specify upper and lower bounds on the influence that one process has on the other persona. Even if unique relationships may exist, the agents may lack the computational ability to analyze a complex body of evidence in order to estimate them. In such situations, reliance on a unique relationship is unwarranted and unwise, and it would be prudent to represent the relationships between processes by a set of social utilities. In such cases, the influence relationships are *imprecise*. Furthermore, it may also be true that the models for θ will not be precise.

An important sub-discipline of probability theory is so-called *imprecise probability theory* [31], which is concerned with making statistical inferences that involve sets of probability measures rather than the classical approach involving a single measure. This theory may serve as the foundation for the development of *imprecise praxeology*, where the social utilities (selectability and

rejectability) are imprecise. This approach represents a natural extension of basic (precise) satisfying theory.

Graph theory also provides a convenient language with which to characterize imprecise systems. A set of probability measures is called a *credal set* [14]. To distinguish between the classical epistemic application and the praxeic application, we will refer to sets of conditional influence relationships as *praxeic sets*. Let V_i be an arbitrary vertex of a DAG, and let $P_{V_i|pa(V_i)}$ denote a set of mass functions of the form $p_{V_i|pa(V_i)}$, with the interpretation that each such conditional mass function is a valid candidate as the true characterization of the relationship between the vertex and its parents (here, the parents may be both praxeic or epistemic vertices, or both). Typically, the set $P_{V_i|pa(V_i)}$ will be taken to be convex; that is, if $p_{V_i|pa(V_i)} \in P_{V_i|pa(V_i)}$ and $p'_{V_i|pa(V_i)} \in P_{V_i|pa(V_i)}$, then all convex combinations of these two mass functions are also in the credal set; that is, for any $\alpha \in (0, 1)$, we have

$$\alpha p_{V_i|pa(V_i)} + (1 - \alpha) p'_{V_i|pa(V_i)} \in P_{V_i|pa(V_i)}.$$

By viewing the influence relationships between vertices as praxeic sets, the praxeic network may be extended to include set-valued probability mass functions. The resulting set of joint mass functions for the entire network is then of the form

$$P_{V_1 \dots V_{2n+m}} = \left\{ \prod_i p_{V_i|pa(V_i)} : p_{V_i|pa(V_i)} \in P_{V_i|pa(V_i)} \right\}.$$

3.3 Key Research Issues

- A unique feature of this approach to multiagent system modeling in the presence of uncertainty and ignorance is that it proposes to deal with both concepts within the same mathematical structure, namely, graph theory. Furthermore, this approach combines both praxeic and epistemic issues into the same unified approach. However, extended study of such networks has never been undertaken, and there is much to learn about the potential of such an approach. To assess the performance of this methodology it will be necessary to develop simulations of potential applications of real-world military interest, including such potential applications as reconnaissance and surveillance, team formation, intelligent guidance and control, and resource allocation.
- The study of credal networks is a current research topic for several theorists [6, 9–11, 29, 32]. One of the major difficulties associated with credal networks is that the number of vertices of credal sets that are combined tends to grow very fast, due to the number of combinations of vertices that must be considered in order to define the credal sets. The current technology in dealing with credal networks can effectively deal with non-densely connected credal networks consisting of approximately two dozen precise vertices; beyond that only approximations are available.

One promising approach, which has been applied to a related problem by Morrell and Stirling [19] is that of set-valued Kalman filtering, where a convex set of probability distributions is propagated via Bayes' rule and dynamic system updates to produce a set-valued estimate rather than the usual point-valued estimate. The key to that approach is to parameterize the credal sets and propagate the parameters, rather than the entire set of estimates. This approach has led to a tractable solution. A potentially fruitful avenue of research would be to apply the same concept to imprecise praxeic networks.

4 An Information-Theoretic Analysis

Satisficing game theory (SGT) provides a coherent mechanism by which agents X_i and X_j can coordinate, but it does not yet indicate the robustness of that coordination to changes in agent's valuations. We propose to develop quantitative measures of agents' commitment to their various options and design methods which exploit these measures. For example, if X_i and X_j can find an option u that is both jointly and individually satisficing *and* to which both are strongly committed, one can expect the resulting coordination between them to be robust. If, on the other hand, the commitment of either agent toward option u is small, coordination may be intermittent or otherwise unreliable. Similar considerations apply to negotiation and group formation. If mutually satisficing solutions cannot be found and the commitment of the various agents to their existing individually satisficing sets is simultaneously large, this may signal that further negotiations will not be fruitful. Better in these circumstances to either change strategies, abandon negotiations, or in the case of group formation, to exclude certain agents from further consideration.

To develop such measures, we propose an information-theoretic approach. Information theory (IT) makes use of scalar measures such as entropy, the Kullback-Leibler measure, and mutual information to characterize the relationships between the probability mass functions (pmfs) which represent the input and output signals of a communications channel. By doing so, the effects of channel noise and interference on communication performance can be obtained. We propose to use such measures to study to what degree agents can modify their own and others' decisions, that is, the content of their satisficing sets, by changing their selectability and rejectability utilities within the constraints imposed by the overall system. Unlike IT, where the signal and channel pmfs have only an indirect effect on the design of error correcting codes, the selectability and rejectability pmfs directly determine the satisficing set, and we expect these methods to apply immediately to the design of SGT systems.

4.1 Review of Entropy and Mutual Information

The entropy $H(Z)$ of the univariate random variable Z with pmf p_Z is defined as

$$H(Z) = H(p_Z) = - \sum_{i=1}^n p_Z(i) \log p_Z(i). \quad (1)$$

In IT entropy is interpreted as the prior uncertainty of the value taken on by Z . Alternatively, it is the average information, or reduction of prior uncertainty, produced by the observation of Z . The mutual information between two random variables Z and W is defined as

$$I(Z; W) = I(p_Z; p_W) = H(Z) - H(Z|W) = H(W) - H(W|Z). \quad (2)$$

$I(Z; W)$ is the average amount of information, or the reduction of prior uncertainty, one can obtain about the value taken on by one random variable by observing a different random variable. $I(W; Z)$ is a special case of the Kullback-Leibler measure $D(p_{WZ}, p_W p_Z)$, which in this case measures the 'distance' between the joint pmf p_{WZ} and the product of marginals $p_W p_Z$, or equivalently, between the conditional pmf $p_{W|Z}$ and the marginal p_W .

The channel capacity problem of IT is to determine for a given channel $p_{Z|W}$ the input distribution p_W that maximizes $I(Z; W)$. Channel capacity defines the highest rate at which information

may be transmitted across a channel with arbitrarily high reliability. The rate-distortion problem is to determine for a fixed p_W the quantizer $p_{Z|W}$ which minimizes $I(Z; W)$, subject to a constraint on the accuracy with which Z represents W .

4.2 Application to Multi-agent Decision Theory

In the praxeological context $H(p_S)$ is interpreted as the *indifference* of the selectability process S , and $H(p_R)$ as the *indifference* of the rejectability process R , to their options. Similarly, the analogue to mutual information is *mutual influence*. Thus, for example,

$$I(S_i; S_j) = H(S_i) - H(S_i|S_j) = H(S_j) - H(S_j|S_i) \quad (3)$$

is the change in S_i 's *a priori* indifference caused by consideration of S_j 's indifference. As with mutual information, mutual influence is symmetric: each agent influences the other equally, and conditioning never increases average indifference. Identical conclusions hold for $I(R_i; R_j)$, $I(S_i; R_j)$, and $I(R_i; S_j)$.

The satisficing set Σ_q^i is determined by the likelihood ratio $\Lambda_i = p_{S_i}/qp_{R_i}$. For a fixed q and p_{R_i} , there is a region $\mathcal{R}_q^i(u)$ of the probability space such that option $u \in \Sigma_q^i$ if and only if $p_{S_i} \in \mathcal{R}_q^i(u)$. The satisficing region $\mathcal{R}_q^i$ associated with the satisficing set Σ_q^i is defined by

$$\mathcal{R}_q^i = \bigcap_{u \in \Sigma_q^i} \mathcal{R}_q^i(u). \quad (4)$$

(The satisficing region may also be defined in terms of q and p_{S_i} .) Because, speaking geometrically, p_{S_i} is not generally equidistant from all boundaries of $\mathcal{R}_q^i$, the options $u \in \Sigma_q^i$ are not equally susceptible to modification by adjustments in q or by the influence of other agents. Define the selectability *commitment* of agent i at level q to option u as

$$C_q^i(u) = \min_{p^* \notin \mathcal{R}_q^i} D(p_{S_i}, p^*). \quad (5)$$

The selectability commitment measures the minimum change in selectability required for an agent to eliminate option u from its satisficing set. We also define

$$C_q^i(u) = \min_{p^* \in \mathcal{R}_q^i} D(p_{R_i}, p^*). \quad (6)$$

to measure the minimum change in selectability required for an agent to include option u in its satisficing set. Definitions for rejectability commitment are similar. The goal of robust design is to maximize the minimum commitment of the coordinating agents to the desired option u .

Because of the various and complex ways agents may interact, we do not expect this to be a trivial problem. For example, if X is perfectly indifferent to its options, so that both p_S and p_R are uniform, then $C_q(u) = 0 \forall u$. If, on the other hand, p_S and p_R are quite different $C_q(u)$ can be quite large for some, but not necessarily all, satisficing options. Note also that, although large indifference in both selectability and rejectability implies a small commitment, small indifference does not necessarily imply large commitment. Consider for example $p_{SR}(u, v) = 0$ for all but one option, so that $H(p_S) = H(p_R) = 0$ and X is not at all indifferent. If that option lies on

the diagonal, i.e. $p_z SR(u, u) = 1$ for some u , then $p_S(u) = p_R(u)$ for all u and $C_1^i(u) = 0 \forall u$. Thus X can be conflicted by strong but offsetting opinions about the selectability and rejectability of option u and, despite its strong views, be completely indecisive. Although these behaviors are intuitively satisfying, they complicate the design problem.

Now suppose that X_i is not coordinating effectively in some setting, and we wish to eliminate option u from Σ_q^i so as to make X_i more deferential to agent X_j . Consider X_j 's average ability to persuade X_i , that is, to change X_i 's satisficing set and therefore its decisions. For all options for which

$$C_q^i(u) > I(S_i; S_j), \quad (7)$$

S_j 's influence on S_i is insufficient to eliminate u from Σ_q^i . Similar conclusions can be drawn about the influence of R_j on S_i , or about X_j 's ability to have X_i add u to its satisficing set. The design problem here would be to design $p_{S_i|X_j}$ so as to achieve the desired coordination.

4.3 Key Research Issues

Although much remains to be done, further analysis along these lines promises to lead to a better understanding of the limits of coordination and to better design methods. We propose to focus research efforts on the following problems:

- A more thorough and detailed exploration of the relationships between indifference, mutual influence, and commitment, and their generalization to continuous action spaces. A connection to the treatment of decision problems in mathematical statistics is particularly warranted, and the use of alternatives to the Kullback-Leibler measure, such as the Jensen-Shannon measure [15], may allow a compact representation of the simultaneous effects of $S_j R_j$ on Σ_q^i .
- The development of numerical design methods which account for agent commitment. First efforts will be based on modifications of the well-known Arimoto-Blahut algorithm [8] for computing channel capacity and rate-distortion curves.
- Computer simulations to test the effectiveness of these design methods.
- Psychological studies to test the relevance and accuracy of the model of commitment to human decision-making.

5 Psychological Studies

Decision making is a central activity of the social and behavioral sciences, philosophical, and engineering disciplines. As described by Simon [23], these approaches fall generally into two main categories: *substantive rationality* and *procedural rationality*. Substantive rationality is normative, and focuses on *why* decision makers should reach conclusions, and procedural rationality is descriptive and focuses on *how* decision makers should reach conclusions. Typically, substantive rationality is implemented by maximizing expected utility, and is the basis for game theory and optimization-based methodologies. Procedural rationality, on the other hand, is typically implemented by invoking rules and heuristics that have been shown by observation to be valid characterizations of behavior in similar circumstances.

5.1 Satisficing Rationality

Substantive rationality and procedural rationality represent two extreme concepts. The former is an expression of the *superlative degree*, where decision makers seek the globally “best” solution according to an explicit quality measure, and the latter is an expression of the *positive degree*, where decision makers seek a “good” solution without invoking explicit measures of quality. Satisficing decision theory provides a decision methodology that lies between these two extremes, and represents the *comparative degree* in the sense that two utilities are invoked to establish that the attributes that favor selection of an option are “better” than the attributes that favor its rejection.

There is considerable colloquial evidence to support this concept of decision making. People routinely compare the pros versus the cons, the upside versus the downside, the pluses versus the minuses, the benefits versus the costs, etc. This way of making decisions is more primitive than a total rank-ordering of options (superlative), but is more sophisticated than simply following heuristic rules (positive). Although the economic, psychological, philosophical, and engineering/computer science literatures are replete with discussions of the two extreme notions, they are relatively silent regarding this notion of comparative decision making. Yet, the idea of viewing an action from two perspectives — one that focuses on the positive consequences of adopting it and another that focuses on the negative consequences — is a powerful concept, and one for which both psychological and mathematical analyses are long overdue.

Studies in behavioral economics are providing an increasingly compelling body of evidence that (a) humans often are not utility maximizers and (b) people are motivated by social concerns as well as by individual concerns when making decisions [4, 5]. Consequently, the neoclassical approach to decision making, especially in group settings, has been called into question. Much work has been done to incorporate social issues such as fairness, inequity aversion, kindness-reciprocity, and social welfare into standard utility theory. [2, 3, 7, 12, 16] The fact remains, however, that standard utility theory is explicitly designed to facilitate expected utility maximization. Social utilities as defined in the context of SGT, on the other hand, are explicitly designed to accommodate social issues, and therefore may be better than standard utilities for as a vehicle for testing social/psychological attributes.

For example, a potential application of SGT is as a vehicle with which to describe and analyze the interaction between tribal groups. Each tribe may have its own set of cultural values which distinguish it from others, but economic imperatives require the tribes to interact with each other. As they interact, the possibilities for cooperation may prompt them to coordinate their activities for a common greater good, such as forming a centralized government. Or they may choose to remain separate and pursue their own parochial interests. A well-known game-theoretic model for intergroup behavior involves the so-called Stag Hunt game [24], a multiagent game in which the players can choose to cooperate to hunt stag or can defect and hunt hares. Considerable analysis using conventional game theoretic approaches has been devoted to this game. SGT, however, may provide a theoretical setting that is more amenable to modeling social relationships than classical game theory provides.

5.2 Key Research Issues

SGT is a new, mathematically rigorous approach to multiagent decision making. It is increasingly the focus of studies designed to investigate its performance and computational characteristics.

Although these preliminary results are encouraging, there is a strong need to perform detailed psychological studies to test the hypothesis that satisficing decision theory is a valid model of human social behavior.

Two desiderata will guide the research program proposed here:

1. the identification of non-zero-sum games that include: (a) provision for repeated play involving two or more players; (b) a measurable pair of attributes suitable for instantiation in the selectability and rejectability processes of SGT; and (c) the potential for instantiation in a face-to-face situation as well as situations in which players compete against each other via computers or against computers;
 2. a comparison of the performance of SGT against at least one other model vis-a-vis the data from specific games in a way that establishes the models' comparative success in accounting for the data (e.g., outcomes consistent with the predicted satisficing solution sets in the case of SGT).
- The first component of the program will be an extensive and systematic review of the extant literature in behavioral economics, which is a loose consortium of microeconomics, cognitive and decision sciences, behavioral psychology, social psychology, anthropology, political science, philosophy, etc. The review will focus on (1) existing non-zero-sum games that meet the criteria described above and (2) mathematical models suitable for comparison to SGT. In addition to identifying existing games, the review will also identify procedures that could be adapted to become game-like and involve two or more players. Such procedures might well be found in the experimental literatures dealing with attention, memory, pattern recognition, problem solving and decision making, contingency perception, impulse control, cooperation and other prosocial behavior, etc.
 - The second component will consist of the design and implementation of several game protocols derived from the review. The protocols will represent a range of behavioral and cognitive variables, a range of degrees and types of social interaction, and a range of number of players. In each case, the validity of the protocol will be established by a panel of expert consultants who will attest that the protocol is adequate to provide and measure the model-specific variables of interest and thereby supply the comparative analysis of models, including SGT, described above. The final component of the proposed research will also draw from the review. A subset of game protocols will be identified as suitable implementations of the category of social interaction known as negotiation. Subjects in the research that uses these protocols will be carefully selected to provide three types of game scenario: (1) the players all share the same ethnic background, and (2) the players are from different ethnic backgrounds. This design will allow a cross-ethnic (each different ethnic group is independently exposed to the same protocol) as well as multi-ethnic (subjects from different ethnic groups are brought together in the same protocol) analysis of the comparative performances of the models.

References

- [1] K. J. Arrow. *Social Choice and Individual Values*. John Wiley, New York, 1951. 2nd ed. 1963.
- [2] G. Bolton. A comparative model of bargaining. *Amer. Econ. Rev.*, 81:1096–1136, 1991.
- [3] G. E. Bolton and A. Ockenfels. A stress test of fairness measures in models of social utility. *Econ. Theory*, 24(4), 2005.
- [4] C. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton Univ. Press, Princeton, NJ, 2003.
- [5] C. Camerer, G. Lowenstein, and M. Rabin, editors. *Advances in Behavioral Economics*. Princeton Univ. Press, Princeton, NJ, 2004.
- [6] J. Cano, M. Delgado, and S. Moral. Propagation of uncertainty in dependence graphs. In *European Conf. on Symbolic and Quantitative Approaches to Uncertainty*, pages 42–47, 1991.
- [7] G. Charness and M. Rabin. Understanding social preferences with simple tests. *Quart. J. Econ.*, 117:817–869, 2002.
- [8] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley, New York, 1991.
- [9] F. Cozman. Credal networks. *Artificial Intelligence*, 120:199–233, 2000.
- [10] J.C.F. da Rocha and F. G. Cozman. Inference with separately specified sets of probabilities in credal networks. In *Second int. Symp. on Imprecise Probabilities and Their Applications*, pages 112–121, 2002.
- [11] E. Fagioli and M. Zaffalon. An exact interval propagation algorithm for polytrees with binary variables. *Artificial Intelligence*, 106:77–107, 1998.
- [12] E. Fehr and K. Schmidt. A theory of fairness, competition, and cooperation. *Quart. J. Econ.*, 114:817–868, 1999.
- [13] R. M. Hogarth and M. W. Reder, editors. *Rational Choice*. Univ. Chicago Press, Chicago, 1986.
- [14] I. Levi. *The Enterprise of Knowledge*. MIT Press, Cambridge, MA, 1980.
- [15] J. Lin. Divergence measures based on the shannon entropy. *IEEE Trans. Info. Theory*, 37(1):145–151, 1991.
- [16] G. Loewenstein, M. Bazerman, and L. Thompson. Social utility and decision making in interpersonal contexts. *J. Personality and Social Psychology*, 57:426–441, 1989.
- [17] R. D. Luce and H. Raiffa. *Games and Decisions*. John Wiley, New York, 1957.
- [18] J. J. Mansbridge, editor. *Beyond Self-Interest*. Univ. of Chicago Press, Chicago, 1990.

- [19] D. R. Morrell and W. C. Stirling. Set-valued filtering and smoothing. *IEEE Trans. Systems, Man, Cybernet.*, 21(1):184–193, January/February 1991.
- [20] J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, CA, 1988.
- [21] H. Simon. Theories of bounded rationality. In C. McGuire and R. Radner, editors, *Decision and Organization*. North Holland, Amsterdam, 1972.
- [22] H. A. Simon. A behavioral model of rational choice. *Quart. J. Econ.*, 59:99–118, 1955.
- [23] H. A. Simon. Rationality in psychology and economics. In R. M. Hogarth and M. W. Reder, editors, *Rational Choice*. Univ. Chicago Press, Chicago, 1986.
- [24] B. Skyrms. *The Stag Hunt and the Evolution of Social Structure*. Cambridge Univ. Press, Cambridge, UK, 2004.
- [25] M. Slote. *Beyond Optimizing*. Harvard Univ. Press, Cambridge, MA, 1989.
- [26] E. Sober and D. S. Wilson. *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Harvard Univ. Press, Cambridge, MA, 1998.
- [27] W. C. Stirling. *Satisficing Games and Decision Making: with applications to engineering and computer science*. Cambridge Univ. Press, Cambridge, UK, 2003.
- [28] W. C. Stirling, R. L. Frost, and H. L. Miller Jr. A socially coherent rational decision theory. Submitted for publication.
- [29] B. Tessem. Interval probability propagation. *Int. J. of Approximate Reasoning*, 7:95–120, 1992.
- [30] A. Tversky and D. Kahneman. Rational choice and the framing of decisions. In R. M. Hogarth and M. W. Reder, editors, *Rational Choice*. Univ. Chicago Press, Chicago, 1986.
- [31] P. Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London, 1991.
- [32] M. Zaffalon, K. Wesnes, and O. Petrini. Reliable diagnosis of dementia by the naive credal classifier inferred from incomplete cognitive data. *Artificial Intelligence in Medicine*, 29:61–79, 2003.

0225

REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering the required data, reviewing the collected data, completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE Final		3. DATES COVERED (From - To) 07/15/2004 - 07/14/2005	
4. TITLE AND SUBTITLE Coding Theory, Information Theory, and Radar				5a. CONTRACT NUMBER FA9550-04-1-0450	
				5b. GRANT NUMBER F	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Calderbank, A. Robert				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Princeton University Fine Hall / Washington Road Princeton, NJ 08544				8. PERFORMING ORGANIZATION REPORT NUMBER Project/Grant #109-6058 COEUS # 004028-001	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Office of Scientific Research (AFOSR)				10. SPONSOR/MONITOR'S ACRONYM(S) AFOSR	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT N/A Approved for public release distribution unlimited					
13. SUPPLEMENTARY NOTES N/A					
14. ABSTRACT This project provides a theoretical foundation for the development of algorithms to significantly expand radar functionality, given a facility capable of high levels of temporal and spatial diversity of waveforms and polarization, where the mode of operation can be rapidly scheduled. The long term aim is to exploit the diversity and flexibility provided by this advanced functionality through adaptive multidimensional waveform and polarization scheduling based on environment modeling and tracking, as well as multi-dimensional adaptive processing of the returns. In collaboration with Bill Moran (Melbourne) and Stephen Howard (DSTO, Australia) we found that the discrete Heisenberg-Weyl group provides a unifying framework for a number of important sequences significant in the construction of phase coded radar waveforms, in communications as spreading sequences, and in the theory of error correcting codes. Among the sequences which can be associated with the Heisenberg-Weyl group are the Golay or Welti sequences, which are pairs of sequences of unimodular complex numbers such that the sum of their individual auto-correlation functions forms delta spike or thumb tack. Results have been shared with the mainstream signal processing community at ICASSP'05 (Philadelphia) and with the defense signal processing community at DASP'05 (Homestead, Utah).					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON A. Robert Calderbank
a. REPORT UNCLASSIFIED	b. ABSTRACT UNCLASSIFIED	c. THIS PAGE UNCLASSIFIED			19b. TELEPHONE NUMBER (include area code) 609-258-4262

6-15-05