Probabilistic Coalition Formation in Distributed Knowledge Environments

Edward A. Billard and Joseph C. Pasquale

Abstract— **In a distributed system, a group of agents have a potential for improved performance depending on their ability to utilize shared resources. This potential synergy raises the question of whether agents should work together in a system-wide group, i.e., a coalition, or whether they should work alone. In general, there is uncertainty as to whether a coalition will form; this uncertainty can arise for various reasons, such as adaptive strategies of the agents or random faults in the system. In this paper, we present a model for performance based upon the probability of coalition formation. The results indicate a limit in potential performance for adaptive agents and, in particular, the global and local maxima along with regions of non-stability. In addition, the model shows how performance is affected by the knowledge environment of the distributed system, that is, the architecture of the system with respect to the distribution of information. Four environments are examined as illustrations of these general categories of information distribution: global information; inaccessible information; local information residing in autonomous agents; and information residing in a master control agent. The results show the distinctions between the environments with respect to probabilistic coalition formation and also demonstrate the loss in environments without communication as compared to a baseline communication environment.**

I. Introduction

In a distributed system, a group of agents have a potential for improved performance based on their ability to utilize shared resources. These agents may be schedulers which allocate processors for jobs [1] or robots that cooperatively handle tasks [2]. This potential synergy raises the question of whether agents should work together in a system-wide group or whether they should work alone (one of the questions that arises in dynamic group formation [3], [4] or the evolution of cooperation [5]). We present a particular model of *coalition* formation to describe the general concept of "working together" in a group.

The main advantage to working alone is that coordination problems, which can defeat the synergy of the coalition, are eliminated, or at least, minimized. However, agents working alone give up the potential performance improvement or benefit offered by the coalition. In general, there is an uncertainty concerning whether the agents will work together or alone which we model as a probability, the *clustering parameter* c, where

$$c = \Pr[\text{agents will work together}],$$
$$1 - c = \Pr[\text{agents will work alone}].$$

We have chosen a probability of coalition formation for two reasons. First, the uncertain and dynamic environments of most distributed systems argue for adaptive techniques that allow agents to make good decisions under different circumstances. These adaptive techniques are typically a probabilistic mechanism to search for "better" performance, and can concern the choice of co-workers. Second, there is always some non-zero probability of failure in real distributed systems. In the next section, we provide a review of the literature to illustrate both of these points.

In this study, we treat the clustering parameter as an independent variable to show the limits of potential performance that an adaptive agent would experience, in particular, the global and local maxima along with regions of non-stability. We call the general shape of the curve defined by these limits the performance *landscape*. This landscape is affected by the *knowledge environment* of the distributed system, that is, the architecture of the system with respect to the distribution of information. The amount of knowledge in an environment is described by the knowledge function $k(n)$ where

$$k(n) = \Pr[\text{at least } n \text{ agents know a fact}].$$

This simple metric says that the amount of knowledge is proportional to the number of agents n (out of a system of N agents) which know a fact. This probabilistic formulation models uncertainty in a distributed system.

The distribution of information, as determined by the amount and quality of communication along with control relationships, affects the corresponding knowledge function and its probabilistic behavior. For example, we examine distributed environments which we consider to be illustrations of the following four general categories of information distribution: global information; inaccessible information; local information residing in autonomous agents; and finally, information residing in a master control agent. Typically, large amounts of highly reliable communication tend to increase global information, and highly asymmetric control relationships such as master-slave relationships tend to increase the information residing in a particular agent.

In Section III.C, the four knowledge environments are shown to cover a knowledge space. The goal of this study is to model the landscape of potential performance of agents

Manuscript received September 12, 1992; revised February 24, 1993 and March 24, 1994. This paper was supported in part by the National Science Foundation, DEC, IBM and NCR.

E. Billard is with the Faculty of Computer Science and Engineering, University of Aizu, Aizu-Wakamatsu City 965-80, Japan.
J. Pasquale is with the Department of Computer Science and Engineering, University of California, San Diego, La Jolla, CA 92093-0114 USA.

IEEE Log Number 9405715.

in these environments with respect to their search for optimal coalition formation.

The paper is organized as follows: Section II presents related work; Section III describes the model in terms of knowledge, coalitions and decisions; several simple examples are presented in Section IV; Section V shows analysis of optimal knowledge and our conclusions are presented in Section VI.

II. Related Work

In this section, we describe related work in knowledge, group formation, and random faults.

A. Knowledge

The concept of knowledge plays an important part in distributed systems and, in particular, affects the design of distributed protocols and decision-making agents. Communication is used to disseminate information so that processes running a distributed protocol move towards the desired result (i.e. to establish some invariant). In distributed applications involving decentralized decision-making, agents use communication to make good decisions concerning the application domain. The communication of information between agents raises an important problem: at what point can an agent assume that every other agent shares a minimum level of knowledge concerning this information? Communication can be viewed as an attempt to increase the minimum level of knowledge in a system.

The effect of communication on knowledge arises in a very obvious manner in a distributed protocol when an agent sends a message containing a fact. An acknowledgement from the receiver of the message lets the sender "know" that the receiver "knows" the fact. But how does the receiver know that the sender knows this yet? Common knowledge [6] is attained when everyone knows that everyone knows *ad infinitum* a fact, but this knowledge cannot be attained in the most general (e.g. unreliable) distributed systems. For an excellent discussion of knowledge, we refer the reader to [6], who cite studies in philosophy [7], artificial intelligence [8] and psychology [9]. Common knowledge in game theory is of particular interest in economics [10], [11], [12], [13].

The impossibility of achieving common knowledge (unless all agents support the fact simultaneously based onreliable and timely information) leads to other approaches in decision-making [14]. In particular, agents may reason recursively [2] [15] based on probabilistic views of other agents' behavior; the recursion is terminated by an equilibrium or other criterion, thus avoiding the infinite recursion found in common knowledge problems.

We are particularly interested in agents that make decisions but avoid the overhead of communication. As a baseline, we consider a global information environment where communication is guaranteed to establish common knowledge. Our other environments are also extreme in that they are communication-free and together illustrate a knowledge space (see Section III.B). Other research examines intermediate environments where the communication itself is probabilistic, in an attempt to communicate only those messages that improve performance [15].

In this study, the performance of the agents is modeled by payoff matrices; simple illustrations are the Prisoner's Dilemma [5], [16], [17] and a load balancing application [4]. Although we do not elaborate on more complicated tasking and planning environments, these too may be refined into a payoff matrix. The Rational Reasoning System [18] considers the effects of hierarchical planning and time-dependent calculations to establish expected payoffs for a robotic application. In [19], a neural network is trained on a set of rules, representing the knowledge about the cooperation of autonomous agents.

B. Group Formation

Group formation is important in distributed systems as agents may take advantage of the inherent synergy of the group in order to derive a mutual benefit. Adaptation is important so that individual agents may search for other compatible agents, with instances in both living organisms and distributed computing systems. Our approach has been influenced by computational ecologies [20] and by evolution, both in cooperation [5] and game-theory studies [21]. The environments analyzed in these studies are characterized by a high degree of uncertainty and changeability; it is not apparent which agents are the most compatible in terms of resources, constraints and strategies.

In our other studies [3], [4], [22], the probability of a decision is subject to feedback and adaptation; that is, the agents search for optimal coalition formation. In the study we present in this paper, the probability of coalition formation is not subject to feedback, but is an independent variable which allows us to focus on the performance landscape that an adaptive agent experiences in environments with different amounts of knowledge.

C. Random Faults

Besides adaptive mechanisms, group formation in distributed systems can be affected by faults which, from an agent's perspective, occur randomly. Examples of these random faults include communication network breakdown, processor and disk crashes, memory faults, etc. The end result is that an agent in a distributed system faces a possibility of an obstacle to the successful interaction with other agents in the system.

It is to the benefit of the agent to attempt to overcome these obstacles: distributed systems have evolved out of centralized systems, in part, to provide a group benefit of sharing information and resources. However, a good decision in the context of working together with other agents may be a very poor decision in the context of working alone. The agent has an uncertainty about which context is currently appropriate.

We are interested in the abstraction of faults in the related areas of group formation, group interaction, and group coordination. Similarly, game theory is concerned with the abstractions of coordination, cooperation and coalitions [23], [16], [17]. A striking example of coordination failure, and a probabilistic solution, is the Byzantine Generals Problem [24] where malicious/defecting generals send misleading messages concerning a coordinated attack.

III. The Model

We develop our model of a system of agents in three stages: introduction of the clustering parameter c for the coalition model (similar to our definition in [4]), definition of the knowledge function $k(n)$, and the description of the decision-making process (or protocols) of the agents. The protocols are chosen for their simplicity and the fact that they allow a demonstration of coalition formation in the context of the knowledge environments.

A. The Coalition Model

We consider particular systems where performance is different based upon the agents working together or working alone, and where the performance can be quantified as static payoffs in a game, representing an average, or expected, performance over time. However, the systems we model also exhibit uncertainty regarding the actual coalition formation, as modeled by the clustering parameter c. This uncertainty can affect which decisions appear to be the best under the circumstances.

Formally, we define a *coalition* as a group whose members have agreed to work together and who expect a payoff, or resultant benefit in performance, that is based upon their collective actions. Otherwise, a decision-maker may decide to act alone as a *singleton* and receive payoffs based only on its own decisions. For now, we consider a 2-agent, 2-action system; see Section III.C for extensions to larger systems. The payoffs are modeled as a game with a set of matrices containing the payoffs associated with each combination of actions:

$$D = \{D^{12}, D^{21}, D^1, D^2\}$$

where

$$D^{ij} = \left[d^{ij}_{kl}\right], \quad D^j = \left[d^j_k\right], \quad k, l \in \{1, 2\}.$$

The matrices D^{12} and D^{21} represent the payoffs to player 1 and player 2, respectively, when they are members of the coalition. When acting as singletons, the vectors D^1 and D^2 represent the payoffs to players 1 and 2, respectively. For example, $d^{12}_{\alpha_1\alpha_2}$ is the payoff to player 1 when choosing action α_1 while player 2 chooses action α_2, where $\alpha_i \in \{1, 2\}$. As singletons, $d^2_{\alpha_2}$ is the payoff to player 2 when choosing action α_2. The *structure* of a game refers to the size, or dimension, of the game matrix; in a 2-agent system, the structure may either be a 2-player matrix or a 1-player vector. The matrix signifies that a coalition has *action dependency* and the vector signifies that agents working alone have *action independency*.

Three example games are presented in Fig. 1 with the pure equilibria in boldface; that is, an action pair (α_1, α_2) provides no motivation, in terms of payoff, for a player to alter its action assuming that the other player does not do so. In Fig. 1(a), action pair (1, 1) in a coalition returns a payoff of 5 to both players and neither can do better by changing its action. It

$$\begin{array}{ccc} & D^{12}, D^{21} & D^1 \\ & \begin{array}{cc} 1 & 2 \end{array} & \\ \begin{array}{c} 1 \\ 2 \end{array} & \begin{bmatrix} \mathbf{5,5} & 0,0 \\ 0,0 & 0,0 \end{bmatrix} & \begin{bmatrix} 0 \\ \mathbf{b} \end{bmatrix} \\ (D^2)^T & \begin{bmatrix} 0 & \mathbf{b} \end{bmatrix} & \end{array} \qquad \begin{array}{cc} \begin{bmatrix} \mathbf{5,5} & 5,0 \\ 0,0 & 0,0 \end{bmatrix} & \begin{bmatrix} 0 \\ \mathbf{b} \end{bmatrix} \\ \begin{bmatrix} 0 & \mathbf{b} \end{bmatrix} & \end{array} \qquad \begin{array}{cc} \begin{bmatrix} 0,0 & \mathbf{5,5} \\ 0,0 & 0,0 \end{bmatrix} & \begin{bmatrix} 0 \\ \mathbf{b} \end{bmatrix} \\ \begin{bmatrix} 0 & \mathbf{b} \end{bmatrix} & \end{array}$$

(a) (b) (c)

Fig. 1. Example Games for Coalition Knowledge.

is easy to find the equilibrium in the singleton vectors: in all cases, the players should select action 2 to get the best payoff, a variable payoff b. These examples will be considered in more detail in Sections IV and V. See [4] for an example that details the payoffs for a job scheduling application.

We now consider the effects of uncertainty in coalition formation as the game is repeated for many stages. A *probabilistic* coalition is a coalition in which the formation decision is based upon the result of a random device that accepts as input the clustering parameter $c \in [0, 1]$. An average 2-player game, $\overline{D} = \{\overline{D}^{12}, \overline{D}^{21}\}$, is induced by the clustering parameter where

$$\overline{D}^{12} = D^{12} \cdot c + \left[D^1 D^1\right](1 - c),$$
$$\overline{D}^{21} = D^{21} \cdot c + \begin{bmatrix} (D^2)^T \\ (D^2)^T \end{bmatrix}(1 - c).$$

An example computation of an average game induced by the clustering parameter is shown in Fig. 4 of [4]. Consider the example in Fig. 1(a). If $c = 1$ then the average game is just the original 2-player game and the combined payoff, or system gain, at equilibrium is $5 + 5 = 10$. If $c = 0$ then the average game is just the replication of the two vectors (because of action independency) and the combined payoff at equilibrium is $2b$. Intermediate values of c induce different equilibria which we examine in more detail in Section IV.

B. The Knowledge Model

We illustrate the knowledge function $k(n)$ by considering four environments, each an N agent distributed system. The environments are illustrations of these general categories of information distribution: global information; inaccessible information; local information residing in autonomous agents; and information available to master-slave agents.

In the *broadcast* environment, a reliable and timely broadcast after the origination of a fact allows all N agents to know the fact; information is shared globally and the agents have a symmetric control relationship. We say that the amount of knowledge in this environment is K_B and that

$$K_B \Rightarrow k(N) = 1.$$

This knowledge level is equivalent to common knowledge where every agent knows that every agent knows *ad infinitum* a fact, given the necessary assumptions regarding reliability and timeliness of communication. In the next three environments, we assume that communication is not allowed or, at least, is not desirable due to overhead considerations.

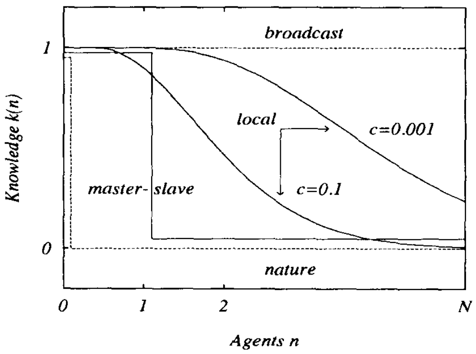

Fig. 2. Probabilistic Knowledge in Four Knowledge Environments.

In the *nature* environment, no agent knows the fact; the information is inaccessible and the agents have a symmetric control relationship. The amount of knowledge in this environment is K_N and

$$K_N \Rightarrow k(1) = 0.$$

In the *master-slave* environment, one agent is guaranteed to know the fact but no other agent knows the fact; the information resides locally in one agent which has an asymmetric control relationship with the other agents. The master has the same level of knowledge as agents in a broadcast environment while the slaves are equivalent to agents in a nature environment. The amount of knowledge in this environment is K_{MS} and

$$K_{\mathrm{MS}} \Rightarrow (k(1) = 1) \wedge (k(2) = 0).$$

In the *local* environment, the information resides locally and each agent has some arbitrary likelihood of knowing the fact. These are autonomous agents with a symmetric control relationship and the amount of knowledge is K_L. We cannot make an implication here but, when the fact concerns coalition formation, we will show the implication in Theorem 1 in Section V.A.

We summarize the four knowledge environments in Fig. 2 where $k(n)$ is plotted for increasing values of n (note that $k(n)$ is actually a discrete function but is drawn continuously for visual purposes). The knowledge in a master-slave environment is a step function between the two extremes of knowledge in the broadcast and nature environments. However, the architecture of a distributed system may generate a knowledge function which is some arbitrary function of n (note that it must be a non-increasing function due to the phrase "at least" n agents know a fact). For example, the knowledge in a local environment is plotted according to the results of Theorem 1 in Section V.A.

It is natural to assume that system performance increases with higher amounts of knowledge; K_B appears better than K_{MS} which appears better than K_N. However, it is not obvious that $k(2) = .5$ (i.e. 50% of the time two agents are certain) is better than $k(1) = 1$ (i.e. the K_{MS} environment where one agent is always certain). In Section V.B, we show

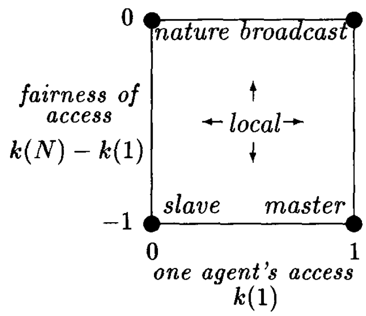

Fig. 3. The Knowledge Space.

analytic results where one environment is preferable to the other environment.

The characteristics of $k(n)$ for each of the environments, reflecting the distribution of information, leads to the two dimensions of a knowledge space shown in Fig. 3. On the x-axis, each agent in a broadcast environment and the master in a master-slave environment has guaranteed access to information. However, each slave in a master-slave environment is unable to access information and, hence, is distinguished from a broadcaster or master. Also, this asymmetry distinguishes a master and a slave from a broadcaster based on the "fairness" of the access to information (the y-axis); that is, *all* of the broadcasters have access to information in their environment but only the master has access to information in its environment. The measure of fairness, $k(N) - k(1)$, is zero in the broadcast environment since the likelihood of all agents knowing the fact is the same as the likelihood of one agent knowing the fact (i.e. the probabilities are one). However, the fairness is -1 in the master-slave environment since $k(N) = 0$ and $k(1) = 1$. At the last corner, the nature environment is as fair as the broadcast environment but, unfortunately, no agent has access to information, $k(1) = k(N) = 0$. Finally, local autonomous agents may populate the entire space.

C. The Decision Model

We have presented models for knowledge environments and coalition formation; in this subsection, we present the decision-making process (or protocols) for agents in different knowledge environments where the fact is "a coalition has formed at this time unit." There are two aspects to the decision-making process. First, each agent acts autonomously and makes probabilistic decisions concerning coalition formation. (These decisions are subject to adaptation in our other studies [3], [4], [22].) The overall likelihood of coalition formation is the clustering parameter c. Second, we assume that decisions concerning actions within a game are made rationally. The agents make the best decision possible under the circumstances of their environment, or equivalently, we are presenting the best-case scenario for adaptive behavior or recursive reasoning [2], [15]. This allows us to concentrate on the uncertainty in coalition formation, however, we believe that adaptation in the action decisions is extremely relevant given the dynamic and uncertain nature of distributed systems.

The clustering parameter c is implemented by each player i flipping a coin to determine whether i is willing to form a coalition with j. The result of a coin flip is $\beta_{ij} \in \{0, 1\}$ and a coalition is formed, for example, in a 2-agent system if, and only if, $\beta_{12} \cdot \beta_{21} = 1$. The coin has an associated probability, the clustering probability c_{ij}, which determines the likelihood that i will agree to coalition formation with j; that is, $c_{ij} = Pr[\beta_{ij} = 1]$ and, in a 2-agent system, $c = c_{12} \cdot c_{21}$. In this study, the c_{ij} are fixed (hence the clustering parameter is fixed) but, in general, can change based on some learning mechanism.

The model does not change significantly for an N-agent system: an N-dimensional matrix D is required for "working together" as a system-wide coalition and this occurs with probability $c = \prod c_{ij}$. (Note that each dimension i of the matrix is indexed by the action of player i, which may have many more choices than just the two shown in our examples.) Otherwise, with probability 1-c, the players act independently and the payoffs are modeled with N vectors (D^i contains the payoffs to player i). Again, an average game $\overline{D}$ can be computed from the large matrix and the individual vectors. Note if any one player decides not to join the coalition, then no intermediate coalitions of smaller size are permitted. This follows one particular game theory treatment [25] where all agents must agree to form the large coalition. The more complex model of every possible coalition of smaller size is not considered here although we have studied 3-player situations in [3].

As a repeated game is played in stages, the coin flips determine the structure of the game at the current stage. However, without a broadcast of information, player i only knows the result of β_{ij}.

We make the following assumptions concerning the common knowledge in the system:

- There is common knowledge of the game matrices.
- There is common knowledge that both players are rational and will select the maximum equilibrium in a game (and a tie-breaking rule is known).
- There is common knowledge of the clustering parameter c and that each player knows its own c_{ij}.

The first two assumptions are only made to allow rational decisions and, hence, eliminate the need to examine adaptive behavior or recursive reasoning in this domain. The last assumption is important since an uncertain, but rational, player must make decisions based upon expected payoffs; without knowledge of c, a player might guess, incorrectly, that each structure is equally likely to occur. A player's knowledge of its own c_{ij} is important since $\prod_j c_{ij} = c$ implies that player i is the master in a *master-slave* environment.

In the broadcast environment, each player i uses communication to broadcast the result of the coin flip β_{ij}; all players know the structure of the game at every stage in a repeated game and the amount of knowledge is K_B. In the master-slave environment, player 1 is invested with a special property: it is the sole source of randomization and, as such, player 1's coin flips exactly determine the structure of the game. Due to the asymmetry, the remaining players never know the structure of

⟵ Knowledge/Uncertainty ⟶

E	K_B broadcast	K_{MS} master-slave (1-2..N)	K_L local	K_N nature
p_i	$\beta_{ij} \leftarrow coin(c_{ij})$ $broadcast(\beta_{ij})$ $receive(\beta_{ji})$ if any $\beta_{ij} = 0$ play D^i else play D	if $i \neq 1$ play $\overline{D}$ else $\beta \leftarrow coin(c)$ if $\beta = 0$ play D^i else play D	$\beta_{ij} \leftarrow coin(c_{ij})$ if any $\beta_{ij} = 0$ play D^i else play $\overline{D}$	$\beta_{ij} \leftarrow coin(c_{ij})$ play $\overline{D}$

Fig. 4. Game Strategy with N Players.

the game but do know the average game based on c. All players commonly know these conditions and we say that the amount of knowledge is K_{MS}. In the nature environment, a global coin c is flipped by nature and all players are uncertain about the structure of the game (K_N). (Alternatively, the players are involved in the randomization process but are unable to interpret the results.) The local environment (K_L) is similar to the nature environment but players test if their own $\beta_{ij} = 0$ and play the singleton vectors if this is the case.

The assumptions concerning the knowledge within an environment lead to the strategies of N rational players presented in Fig. 4. Each player is a synchronous agent making randomized decisions concerning *with whom* to coordinate; the only communication protocol is to broadcast the results of the coin flips in the K_B environment, which then acts as a baseline for the non-communication environments. The amount of knowledge (in environment E) concerning coalition formation increases to the left in the table or, inversely, uncertainty increases to the right. We have placed the local environment, K_L, between the K_{MS} and K_N environments although we show in Section V.B that the performance relation between K_L and K_{MS} depends on the game payoffs.

Besides making randomized decisions concerning *with whom* to coordinate, the agents also make decisions concerning *what action* to perform, which determines game payoffs, but these decisions are made "rationally". Rational players are optimizing agents and since they assume the other player is also rational, they select the best pure equilibrium in the game, that is, the highest combined payoff for an action pair such that either player would suffer if its own action were altered. We assume the games present no difficulties in achieving equilibrium due to competitiveness; that is, the Pareto-efficient action combination is a Nash equilibrium as in coordination games [26]. In addition, we assume a tie-breaking rule is available for multiple Pareto-optima. We admit that this type of simplicity can lead to ad-hoc rules for a game with an equilibrium that is "unfair" to a player (our examples illustrate games without these problems) and more sophisticated reasoning schemes [2] may be candidates for substitution. These assumptions reflect our interest in the effects of the knowledge environment on the players' ability to reach the best equilibrium in the N-player versus 1-player games; we do not focus on the difficulties of rational agents

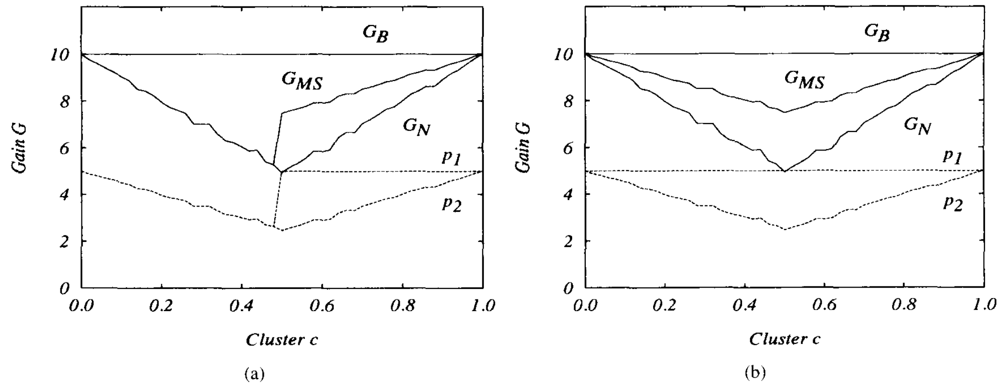

Fig. 5. Effects of Coalition Knowledge (K).

reaching equilibrium in N-player games, which is outside the scope of this paper.

We detail the strategies for the specific case of two players. In the broadcast environment (K_B), both players flip a coin and broadcast the results; the results are received and if any value is zero then it is common knowledge that a coalition cannot form at the current stage. In a coalition, the matrices D^{12} and D^{21} are used for payoffs to player 1 and player 2, respectively. A player *plays* a game by computing the highest equilibrium in $\{D^{12}, D^{21}\}$ and selecting the corresponding action. If a coalition is not formed, i.e. at least one β_{ij} is zero, then each player i plays D^i which is a vector of payoffs and the maximum is chosen.

In the local environment (K_L), both players act the same: if player i flips $\beta_{ij} = 0$ then player i knows, locally, that a coalition cannot form and plays D^i. Otherwise, player i plays $\overline{D}^{ij}$.

In the nature environment (K_N), both players flip a coin (or nature does) and both are uncertain concerning coalition formation. Both players compute the average game $\overline{D}^{ij}$ based on their knowledge of c and play this game by determining the maximum equilibrium.

The master-slave environment (K_{MS}) is more difficult to describe because of the asymmetry: player 1 knows the coalition formation since its coin $c_{12} = c$ whereas player 2 is always willing to form a coalition ($c_{21} = 1$ and $\beta_{21} = 1$). However, player 2 is uncertain about the master's intentions and must play the average game based on its knowledge of c. On the other hand, the master has no uncertainty regarding the appropriate game structure after it flips the coin (these are the differences in knowledge that we attempt to capture in the model). Player 1 examines the result of its own coin flip and if zero plays the vector D^1. However, if a coalition is formed (i.e. D^{12} is the structure), player 1 knows that player 2 is playing the average game $\overline{D}^{21}$ and computes player 2's action α_2 in this game. Now, the column $D^{12}_{\alpha_2}$ is projected and the optimizing player 1 selects the maximum payoff. The resultant action pair may not be an equilibrium but then the other action cannot be one either since player 1 would have selected it. This breakdown occurs because the players are optimizing on different game structures.

Although the coin-flip of the slave is irrelevant, this does not allow the slave any additional freedoms or advantages; the slave must play the "odds" as expressed by the average game since it has no actual knowledge of the true game structure. The master follows the dictates of its coin flip, but it still has complete knowledge of the true game structure. The fact that the decision is made probabilistically has no affect on the masters' level of knowledge, but only on its ability to control the environment. In an adaptive environment, the probability of coalition formation would not be an independent variable and would be under the control of the master as it searches for optimal formation.

IV. Examples of Coalition Knowledge

To illustrate the amounts of knowledge and the strategies of the players, we examine the games in Fig. 1(a) and Fig. 1(b); both games are identical except for one payoff. There is a high degree of symmetry within the games as the maximum payoff is 5 for both players regardless of the game structure since we choose $b = 5$. However, it requires knowledge of the game structure to achieve this payoff since opposite actions are required; the players should select action 1 if a coalition forms, otherwise they should select action 2.

We have chosen simulation to determine the performance of the system because of the difficulties in analyzing the probabilistic payoffs at equilibrium for all four knowledge environments. In the simulations, the probabilities c and $k(n)$ represent a fraction, that is, some number of stages out of the total number of stages in a repeated game.

Fig. 5 shows the performance landscape for two players using the K_B, K_{MS} and K_N strategies presented in Fig. 4 (K_L is considered in the next section). The landscape is defined by the gain G_E, that is, the average combined payoff over time in knowledge environment K_E. The gain is shown as a function of the clustering parameter c for the three environments; each point is the average gain over a run of 2000 stages.

First, we discuss the results in Fig. 5(a), which corresponds to the game in Fig. 1(a). In K_B, both players know exactly, by broadcasting, the status of coalition formation and can choose the appropriate action to receive the payoff of 5 (the landscape is flat because there are identical payoffs in the 2-player game

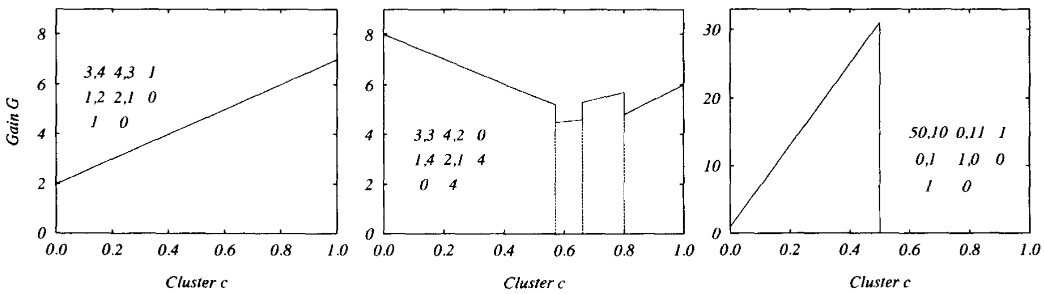

(a) (b) (c)

Fig. 6. Three Types of Equilibria.

and the 1-player games). In K_N, both players compute the average game; for clustering values 0 to .5 it is best to play action 2 since the 1-player game is dominant; for clustering values .5 to 1 it is best to play action 1. When $c = 0$ or 1, these decisions are optimal and agree with K_B; that is, the uncertainty is eliminated at the extrema. However, the players must stick with the same action for half the domain; the changing nature of c decreases the total payoff received and reaches a minimum at $c = .5$.

In the K_{MS} environment, the individual gains of the players are shown in dashed lines (labeled p_1 and p_2). In the lower domain 0 to .5, the gains coincide since the uncertain player is playing action 2, leaving the certain player with a choice; however, the payoff is zero for either choice in the second column of the 2-player game. The knowledge that a coalition is formed does not help player 1 in this domain. Note that both player 1 and player 2 receive a payoff of 5 for whatever fraction of the time they operate as singletons. However, in the higher domain .5 to 1, player 2 chooses action 1 and player 1 always gets a payoff of 5 since it can select action 1 if the coalition forms, or otherwise action 2. Of course, if the independent variable c were under the control of the master, then the master could pick $c = 0$ or $c = 1$ and guarantee that *both* agents receive the maximum of 5. This is exactly the goal of a potential adaptive agent that searches the landscape for optimal coalition formation. However, the payoff at equilibrium in the 2-player game might be greater (e.g. a global maximum) than in the 1-player game (e.g. a local maximum), or vice versa. In another study without distinctive knowledge environments [4], we show that adaptive agents, with sufficient communication, can locate the global maximum. We shall return to this issue shortly.

The K_B segment represents the optimal gain but it also generates the most communication. The area between this segment and the line segments of the lower amounts of knowledge represents a loss due to uncertainty.

Fig. 5(b) has similar results but the payoff is always 5 for player 1 in the K_{MS} environment. Note that the associated game has a payoff of 5 in the second column (i.e. the column selected by player 2 in the lower domain of c); when player 1 detects that a coalition is formed it plays action 1, otherwise action 2. This plot shows an increase in the gain for all clustering values as the level of knowledge increases (except at the extrema where they coincide).

Different game payoffs can generate landscapes with qualitatively different characteristics. For instance, Fig. 6 shows an analytic determination of the gains at equilibrium versus the clustering parameter for three games in a K_N environment. The equilibria are restricted to those in *pure* strategies rather than *mixed* (i.e. probabilistic) strategies for reasons stated below. The plots show that particular values of c are local or global maxima. The nature of the equilibria has implications for agents that search for optimal coalitions within a distributed system by modifying their c_{ij}, hence c, over time. For example, the agents may adapt to a local maximum, at the loss of finding the global maximum, or may operate in a region without a pure equilibrium and, hence, fail to locate an "uphill" direction.

The three plots in Fig. 6 illustrate three categories of *stability*, that is, the extent of the equilibria with respect to the domain of the clustering parameter. Note that all games at least have a pure strategy equilibrium at $c = 0$ since the singleton vectors always provide such an equilibrium. For this reason, we call each category *stable* but with possible qualifiers based upon the extent over the clustering parameter.

In Fig. 6(a), the dominant strategies for the 2-player game coincide with those of the 1-player games; we call this a *strongly stable* equilibrium since the same action pair is in equilibrium over the entire domain of the clustering parameter, hence uncertainty is eliminated. In Fig. 6(b), there are four different equilibria depending on the clustering parameter; we call these *stable* equilibria since there is always some equilibrium available over the domain of the clustering parameter. Note that a global maximum exists at $c = 0$ and that a local maximum exists at $c = 1$ but the discontinuities in the graph imply that rational or adaptive agents must alter action strategies to achieve equilibrium as the clustering parameter changes.

In Fig. 6(c), an equilibrium does not exist, in pure strategies, at high values of the clustering parameter; this game has a *weakly stable* equilibrium at low values. Note that a global maximum occurs at $c = .5$. Of course, the players will get some payoff at high values of c, but these cannot be predicted on the basis of an equilibrium in pure strategies. The analytic treatment makes no prediction concerning the behavior of players under the condition that a pure strategy equilibrium does not exist. Of course, all nonzero-sum games have an equilibrium in, at least, mixed strategies, assuming that both players are rational and know the payoffs to each other [23]. We have observed that adaptive agents, without the presence of a pure strategy equilibrium, slowly cycle through all action pairs as each new potential equilibrium fails with respect to another potential equilibrium [3].

V. Analysis

In this section, we analyze optimality and ordering in the knowledge environments.

A. Optimal Conditions for Knowledge Function

We have described a broadcast environment where both players always know the structure of the game. However, if players only have the local information available in a K_L environment, what is the best choice of c_{ij}, for a given c, such that $k(2)$ is optimized? We consider the local environment to be of particular interest since communication is not required and the agents make autonomous decisions based on local variables. Assuming local information, Theorem 1 shows that $c_{ij} = \sqrt{c}$ maximizes $k(2)$, for example; in general, $c_{ij} = \sqrt[n]{c}$ maximizes $k(n)$.

The broadcast and local environments have equivalent amounts of knowledge when all $c_{ij} = 0$, as each player is aware that a coalition cannot form. Although not indicated by local coin flips, rational agents can deduce that $c = 1$ implies that a coalition will always form. These are the only conditions that remove uncertainty and the players cannot achieve, using only local information, a K_B amount of knowledge for arbitrary values of the clustering parameter c.

Theorem 1 (Optimal Knowledge): *Let $1 \le i, j \le N$, $i \ne j$, and let $c_{ij} \in [0, 1]$ be the probability that player i is willing to form a coalition with j. Let $\beta_{ij} \in \{0, 1\}$ be the result of the randomization at a stage in a repeated game. We say that a coalition is formed at a stage if, and only if, $\prod \beta_{ij} = 1$. Let $c = \prod c_{ij}$ be the probability that a coalition forms at a stage and let $k(n)$ be the probability that at least $n \in [0, N]$ players know that a coalition has formed. For a given c and the fact that player i uses only local information of β_{ij}, $k(n)$ is optimized by*

$$k^*(n) = \begin{cases} 1 & \text{if } n = 0 \\ 1 & \text{if } n = 1, \quad c_{12} = c \\ (1 - \sqrt[n]{c})^n & \text{if } n > 1, \quad \prod_j c_{ij} = \sqrt[n]{c}, \quad i \le n \\ & \qquad\qquad c_{ij} = 1, \quad i > n \end{cases}$$

Proof: By cases,

Case $n = 0$: $k^*(0) = 1$, trivially, since there is a probability of one that at least zero players know that a coalition has formed.

Case $n = 1$: if $c_{21} = 1$, then $c_{12} = c$ and player 1's coin flip, i.e. local information, completely determines the coalition formation and $k^*(1) = 1$.

Case $n = 2$, $N = 2$: with probability $1 - c_{ij}$, player i knows, based on local information, that a coalition cannot form, hence the probability that *both* players know is optimized by

$$k^*(2) = \max_{c_{ij}} (1 - c_{12}) \cdot (1 - c_{21}).$$

Let $f(c_{12}) = (1 - c_{12}) \cdot (1 - c_{21}) = (1 - c_{12}) \cdot (1 - \frac{c}{c_{12}})$, since $c = c_{12}c_{21}$. Then

$$\frac{df(c_{12})}{dc_{12}} = -1 + \frac{c}{(c_{12})^2} = 0$$

$$\text{when} \qquad c_{12} = \sqrt{c} \quad \text{and} \quad c_{21} = \sqrt{c}$$

and the optimum is

$$k^*(2) = (1 - \sqrt{c}) \cdot (1 - \sqrt{c}) = (1 - \sqrt{c})^2.$$

Note that the critical point is a maximum since the second derivative, evaluated at $c_{12} = \sqrt{c}$, is negative:

$$\frac{d^2 f(c_{12})}{dc_{12}^2} = \frac{-2c}{(c_{12})^3} = \frac{-2}{\sqrt{c}} < 0.$$

Case $n = 2$, $N = 3$: without loss of generality, let us maximize the probability that the first two players know that a coalition cannot form. For example, the probability that player 1 knows that a coalition cannot form is $1 - c_{12}c_{13}$ since any $\beta_{1j} = 0$ determines this fact for player 1. Therefore,

$$k^*(2) = \max_{c_{ij}} (1 - c_{12}c_{13})(1 - c_{21}c_{23}).$$

This has the same solution as before except $c_{12}c_{13} = c_{21}c_{23} = \sqrt{c}$ rather than $c_{ij} = \sqrt{c}$. The symmetry is required; for example, if $c_{12}c_{13} > c_{21}c_{23}$, then $k(1)$ increases but $k(2)$ decreases. Note that $c_{31} = c_{32} = 1$ in order to allow the minimum c_{ij}'s for the first two players (i.e. this allows the first two players to know more often).

Case $n > 1$, $N \ge n$: the probability that player 1 knows that a coalition cannot form is $1 - \prod c_{1j}$ and the optimal probability that the first n players *all* know the fact is

$$k^*(n) = \max_{c_{ij}} \prod_i \left(1 - \prod_j c_{ij}\right).$$

Again, there must be a symmetric distribution of c among the first n players, that is, $\prod_j c_{ij} = \sqrt[n]{c}$ for $i \le n$ or else $k(n)$ decreases. As before, $c_{ij} = 1$ for $i > n$. With these conditions in a local environment,

$$K_L \Rightarrow k^*(n) = (1 - \sqrt[n]{c})^n. \qquad \square$$

Theorem 1 also makes additional implications as shown in Fig. 2 for the local environment (the data is $k^*(n)$ for $c = 0.1$ and $c = 0.001$). First, as the value of c decreases, the optimal knowledge increases as more agents know, locally, that a coalition will not form. This does not imply that the performance of the system necessarily improves since the payoffs for working alone may be less than for working together. Second, as n increases, the optimal knowledge tends to zero. This also implies that the optimal probability that all N agents know the coalition fact tends to zero for large systems. The distribution of randomization to a large number of agents diminishes the local knowledge of each agent concerning coalition formation. This result is formalized in the following corollary.

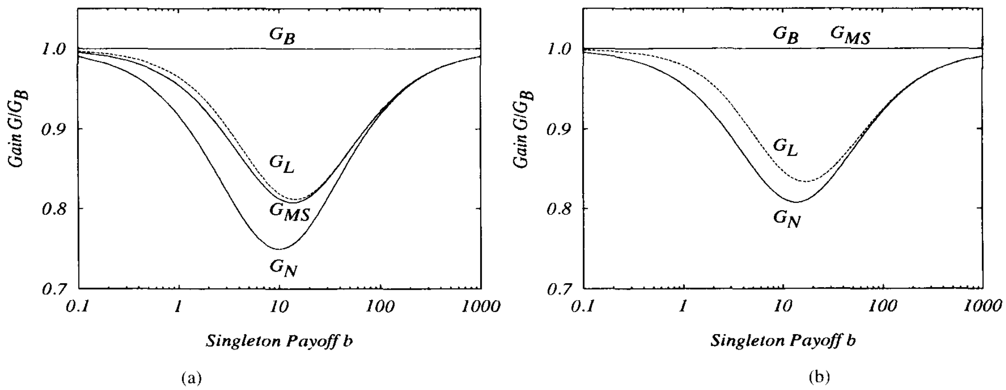

Fig. 7. Gain in Knowledge Environments as a Fraction of G_B.

Corollary 1 (Infinite Agents) For $c > 0$,

$$\lim_{N\to\infty} k^*(N) = \lim_{N\to\infty}(1 - \sqrt[N]{c})^N = 0 \Rightarrow \lim_{N\to\infty} c^*_{ij} = 1.$$

B. Knowledge Ordering

If $\overline{G}_E$ is the average gain in environment E over all c, then we conjecture that, for all games,

$\overline{G}_B \geq \overline{G}_L \geq \overline{G}_N$ since, by definition, $K_B \supseteq K_L \supseteq K_N$

and

$\overline{G}_B \geq \overline{G}_{\mathrm{MS}} \geq \overline{G}_N$ since, by definition, $K_B \supseteq K_{\mathrm{MS}} \supseteq K_N$.

An interesting question is the relation between $\overline{G}_{\mathrm{MS}}$ and $\overline{G}_L$: is it better for one agent to always know or for two agents to sometimes know?

The simulation of performance in different knowledge environments can be replaced with analysis using arbitrary game matrices constrained to equilibrium solutions, however, this does not lead to simple closed-form solutions. Instead, we examine the two extreme game structures in Fig. 1(a) and Fig. 1(c) to illustrate that master-slave and local environments have a fixed ordering within a game type but this ordering may be reversed in a different game type. Fig. 1(a) represents a worst-case scenario for the uncertain player since opposite actions are required, depending on the coalition decision, and a nonoptimal action decision yields a payoff of zero. The game structure does not present any difficulty due to competition between the players; it is only the uncertainty in coalition formation that hinders the players.

We present an analytic comparison of the gains in the different environments in Fig. 7(a), corresponding to the game in Fig. 1(a), over a range of values for b, the singleton payoff (see [3] for a complete derivation). The gain, as a fraction of the G_B baseline, shows that the local environment is better than, or at least as good as, the master-slave environment for this type of game structure. Fig. 7(b) shows the gains that correspond to the game in Fig. 1(c), and that the master-slave environment is better than the local environment. In this best-case master-slave game, the performance in both the master-slave and broadcast environments is identical since the uncertainty is removed for the slave: it always chooses the second action.

VI. Conclusion

We have focused on the problem of agents working together in groups or working alone to demonstrate the effects on performance due to the knowledge environment and the uncertainty in group formation. The landscape of potential performance allows us make several observations. The landscape may have local and global maxima as well as regions where non-stable behavior can be expected. This has implications for agents that search for optimal group formation. For example, an agent may adapt to a local maximum, at the loss of finding the global maximum, or may operate in a region without a pure equilibrium, hence failing to locate an "uphill" direction. There are situations where it is better for one agent to *always* be correct than for two agents to *sometimes* be correct. In a local environment with a large number of agents, the use of randomization (and the non-use of communication) makes the "knowledge" tend toward zero. The graphical representation is a simple illustration of the loss in non-communication environments as compared to a baseline communication environment.

Although the environments present extreme, and sometimes simplified, conditions, our work attempts to model key points in adaptive group formation. We have some evidence [22] that the use of adaptive strategies for group formation improves performance of decentralized job schedulers in a queueing system.

Acknowledgment

The comments from the reviewers and the editorial suggestions from Alice Riedmiller are gratefully acknowledged.

References

[1] T. Casavant and J. Kuhl, "A formal model of distributed decision-making and its application to distributed load balancing," in *Proc. IEEE Int. Conf. Distr. Comput. Syst.*, pp. 232–239, Sept. 1986.

[2] P. Gmytrasiewicz, E. Durfee and D. Wehe, "A decision-theoretic approach to coordinating multiagent interactions," *Proc. 12th Intl. Joint Conf. on AI*, pp. 166–172, 1991.

[3] E. Billard, "Delayed Communication and Dynamic Group Formation in Distributed Systems," *Ph.D Thesis*, University of California, San Diego, 1992.

[4] E. Billard and J. Pasquale, "Effects of delayed communication in dynamic group formation," *IEEE Trans. Syst., Man, Cybern.*, vol. 23, no. 5, pp. 1265–1275, Sept./Oct., 1993.
[5] R. Axelrod and W. Hamilton, "The evolution of cooperation," *Science*, vol. 211, pp. 1390–1396, Mar. 1981.
[6] J. Halpern and Y. Moses, "Knowledge and common knowledge in a distributed environment," *J. ACM*, vol. 37, pp. 549–587, July 1990.
[7] J. Hintikka, *Knowledge and Belief.* Ithaca, NY: Cornell University Press, 1962.
[8] R. Moore, "A formal theory of knowledge and action," in *Formal Theories of the Commonsense World* (J. Hobbs and R. Moore, eds.), pp. 319–358, Norwood, NJ: Ablex Publishing Corp., 1985.
[9] H. Clark and C. Marshall, "Definite reference and mutual knowledge," in *Elements of Discourse Understanding*, (A. Joshi, B. Webber and I. Sag, eds.), pp. 10–63, Cambridge, MA: Cambridge University Press, 1981.
[10] J. Geanakoplos, "Common knowledge," *J. Economic Perspectives*, vol. 6, no. 4, pp. 53–82, 1992.
[11] A. Brandenburger, "Knowledge and equilibrium in games," *J. Economic Perspectives*, vol. 6, no. 4, pp. 83–101, 1992.
[12] P. Reny, "Rationality in extensive-form games," *J. Economic Perspectives*, vol. 6, no. 4, pp. 103–118, 1992.
[13] R. Aumann, "Agreeing to disagree," *Annals of Statistics*, vol. 4, no. 6, pp. 1236–1239, 1976.
[14] J. Rosenschein and J. Breese, "Communication-free interactions among rational agents: A probabilistic approach," in *Distributed Artificial Intelligence* (L. Gasser and M. Huhns, eds.), pp. 99–118, London: Pitman, 1989.
[15] P. Gmytrasiewicz, E. Durfee and D. Wehe, "The utility of communication in coordinating intelligent agents," *Proc. 9th Natl. Conf. on AI*, pp. 166–172, July 1991.
[16] H. Hamburger, "N-person prisoner's dilemma," *J. Math. Sociology*, vol. 3, pp. 27–48, 1973.
[17] W. Poundstone, *Prisoner's Dilemma.* New York: Doubleday, 1992.
[18] P. Gmytrasiewicz, E. Durfee and D. Wehe, "Combining decision theory and hierarchical planning for a time-dependent robotic application," *Proc. 7th IEEE Conf. on AI Applications*, pp. 282–288, Feb. 1991.
[19] E. Szczerbicki, "Acquisition of knowledge for autonomous cooperating agents," *IEEE Trans. Syst., Man, Cybern.*, vol. 23, no. 5, pp. 1302–1315, Sept./Oct., 1993.
[20] B. Huberman and T. Hogg, "The behavior of computational ecologies," in *The Ecology of Computation* (B. Huberman, ed.), North-Holland: Elsevier Science Publishers, 1988.
[21] J. Maynard-Smith, *Evolution and the Theory of Games.* Cambridge University Press, 1982.
[22] E. Billard and J. Pasquale, "Dynamic scope of control in decentralized job scheduling," in *Proc. IEEE Int. Symp. Autonomous Decentralized Syst.*, pp. 183–189, Mar. 1993.
[23] A. Rapoport, *N-Person Game Theory*, Ann Arbor, MI: University of Michigan Press, 1970.
[24] L. Lamport, R. Shostak and M. Pease, "The byzantine generals problem," *ACM Trans. Prog. Lang. and Syst.*, vol. 4, pp. 382–401, July 1982.
[25] N. Howard, *Paradoxes of Rationality: Theory of Metagames and Political Behavior.* Cambridge, Mass.: MIT Press, 1971.
[26] V. Crawford and H. Haller, "Learning how to cooperate: optimal play in repeated coordination games," *Econometrica*, pp. 571–595, 1990.

Edward Billard received the B.S. degree in engineering physics with special honors in 1973 and the M.S. degree in computer science in 1980 from the University of Colorado, Boulder. He received the Ph.D. degree in computer science from the University of California, San Diego in 1992 and continued until 1993 as a postdoctoral scholar in the Computer Systems Laboratory.

Besides several positions in software engineering and management, Dr. Billard has taught physics in the Peace Corps and computer science at the University of Colorado and West Coast University, San Diego. He is currently an Assistant Professor of Computer Science at the University of Aizu, Japan, where he directs the Operating Systems Laboratory. His current research interests include autonomous decentralized systems, adaptation, dynamic group formation, decision making under delayed communication, dynamical systems, queueing systems, and game theory.

Joseph Pasquale is an Associate Professor of Computer Science and Engineering at the University of California at San Diego, where he co-directs the UCSD Computer Systems Laboratory. He is also a Senior Fellow at the San Diego Supercomputer Center.

Dr. Pasquale received the B.S. and M.S. degrees from the Massachusetts Institute of Technology in 1982, and the Ph.D. degree from the University of California at Berkeley in 1988, all in computer science. He is a recipient of the Presidential Young Investigator Award from the National Science Foundation, and faculty development awards from DEC, IBM, NCR, and TRW.

Dr. Pasquale's research interests include operating system and network design, especially to support I/O intensive applications such as distributed multimedia (digital video and audio) and scientific computing. He is also interested in issues of coordination and decentralized control in large distributed systems. He is a principal architect of the Sequoia 2000 network, a wide-area high-speed network spanning the University of California campuses.