

DISTRIBUTED GAMES

DOV MONDERER AND MOSHE TENNENHOLTZ

Faculty of Industrial Engineering and Management,
Technion–Israel Institute of
Technology, Haifa , Israel
e–mail:
dov@ie.technion.ac.il
moshet@ie.technion.ac.il

August 1997

Abstract.

The Internet exhibits forms of interactions which are not captured by existing models in economics, artificial intelligence and game theory. New models are needed to deal with these multi-agent interactions. In this paper we present a new model – distributed games. In such a model each player controls a number of agents which participate in asynchronous parallel multi-agent interactions (games). The agents jointly and strategically control the level of information monitoring by broadcasting messages. As an application, we show that the cooperative outcome of the Prisoner’s Dilemma game can be obtained in equilibrium in such a setting.

Keywords. Distributed Games, Internet.

1. Introduction. The Internet introduces new challenges in artificial intelligence, economics and game theory. It exhibits both parallel and sequential interactions. While sequential interactions have been extensively discussed in the literature, the study of parallel interactions has been neglected so far. New models of economies and games are needed in order to effectively deal with parallel interactions. In this paper we present one such new model – distributed games. Our model captures several features of distributed systems, such as the Internet. Such systems

First Version: April 1997. This work has been influenced by discussions made by both authors with the participants of the artificial intelligence (AAAI) Spring Symposium, Stanford University 1997. We thank Avraham Neyman, Haim Reisman, Ariel Rubinstein, and Doron Sonsino for helpful discussions and suggestions.

are usually assumed to be asynchronous systems where agents communicate by broadcasting messages (Tanenbaum 1988). To motivate our particular definition, think of several users (players) who send software agents¹ to participate in auctions which are held in various locations of the net, and are conducted in a randomly chosen order. This random order of play models a situation where all auctions are supposed to take place in parallel but, due to the fact that real distributed systems are asynchronous, are actually held at different time intervals. It may also reflect a random decision of the auctions' organizers. The agents can communicate by broadcasting messages. These messages are not considered to be private information of the agents². A main feature in our model is that the level of information monitoring³, which is determined by the messages is not exogenous but rather is the result of a joint strategic decision of the players. In a general distributed game the agents may be engaged in any type of interaction at each of the locations. As an application, we analyze a distributed game in which a Prisoner's Dilemma game is played at several locations, and the probability distribution over the order of play is the uniform one. We show that if the number of locations is sufficiently large, the outcome of cooperation is naturally obtained in equilibrium. In this equilibrium, an agent cooperates if and only if it has not gotten a message. If it detects a deviation of its opponent it immediately broadcasts an alarm message, otherwise it does not send any message. Note that this equilibrium survives due to the joint decision of agents not to reveal to their partners and opponents any information regarding the number of future interactions, as long as their opponents cooperate with them. Hence, in equilibrium the agents never know how many more games are to be played. This lack of knowledge enables cooperation.

2. Distributed Games—A Model. We consider m players, $T = \{1, 2, \dots, m\}$, and n m -player games in strategic form, G_i , $1 \leq i \leq n$. The players in G_i are

¹Roughly speaking, software agents are programs which are designed to serve the goals of a particular user. These agents may navigate in a computerized network, while transmitting messages among themselves, and interacting with other agents (which might be controlled by other users). The design of software agents is one of the most important research and application directions of the computer industry CACM (1994).

²This assumption reflects the current state of Internet security. It motivates cryptography theory (e.g., Diffie and Hellman (1976)). In the model of distributed games considered in this paper we do not distinguish between a message (e.g., a stream of bits) and its content. In more advanced models one may wish to deal with asymmetric information setup, where only the agents of the same player can decrypt their messages. The result proved in Section 3 holds for such models too. information.

³that is, the information about the games that have already been played and about the actions that were chosen in these games.

denoted by $1i, 2i, \dots, mi$, where ji is referred to as the agent of Player j at the location i . The action set of Agent ji at the game G_i is denoted by S_i^j . The payoff function of Agent ji at G_i is $u_i^j : S_i \rightarrow R$, where $S_i = \times_{j=1}^m S_i^j$ and R denotes the set of real numbers. Let λ be a fixed probability distribution over the $n!$ random orders of the set of games $\{G_1, G_2, \dots, G_n\}$. Let M^j be a finite set of messages that can be sent by the agents of player j . We assume that all players have the same message set, denoted by M . That is $M^j = M$ for every $j \in T$. It is further assumed that M contains the no-message option, denoted by m_0 . Thus for every Player j , $M^j = \{m_0, \dots, m_k\}$, where k is a non-negative integer. Let M^T be the set of message profiles that can be broadcast from a location. For a location i let $H_{-i} = (M^T)^{\{1,2,\dots,n\} \setminus \{i\}}$ be the set of message profiles that can be transmitted before the agents at location i are called to choose their actions. That is, each $h \in H_{-i}$ is a vector $h = (h_l)_{l \neq i}$, $1 \leq l \leq n$, with $h_l = (h_l^1, h_l^2, \dots, h_l^m) \in M^T$. That is h_l^j is the message sent by the agent of Player j at location l . If $h_l^j = m_0$, it means that Agent jl has not sent a message. Note that if $h_l^j = m_0$ for every $j \in T$ then from the point of view of the agents in location i , either the game at location l has not been played yet, or it has been played but neither of the agents at location l sent a message. A strategy for Player j is a pair $\sigma^j = (f^j, g^j)$, where $f^j = (f_i^j)_{i=1}^n$ and $g^j = (g_i^j)_{i=1}^n$, where $f_i^j : H_{-i} \rightarrow S_i^j$ and $g_i^j : H_{-i} \times S_i \rightarrow M$. That is, for every history of messages $h \in H_{-i}$, Agent ji chooses the action $f_i^j(h)$, and after observing all other agents' actions at the location i , $x_i \in S_i$, it broadcasts the message $g_i^j(h, x_i)$. We assume that the broadcasts are received by all agents in all locations⁴. Every tuple of strategies $\sigma = (\sigma^j)_{j=1}^n$ defines a probability distribution over $S = S_1 \times S_2 \times \dots \times S_n$, denoted by μ_σ . The expectation operator with respect to σ is denoted by E_σ . The expected payoff of agent ji given σ , is $E_i^j(\sigma) = E_\sigma(u_i^j(\cdot))$, where $u_i^j(x_1, x_2, \dots, x_n)$ is identified with $u_i^j(x_i)$. We denote the expected payoff of Player j by $u^j(\sigma)$. That is,

$$u^j(\sigma) = \sum_{i=1}^n E_i^j(\sigma).$$

Let Σ^j be the set of strategies of Player j , then the above mentioned distributed game defines a game in strategic form with the player set $T = \{1, 2, \dots, m\}$. In this game Player j chooses $\sigma^j \in \Sigma^j$ and receives the payoff $u^j(\sigma)$. Equilibrium concepts in the distributed game are inherited from the equilibrium concepts in the associated game in strategic form.

⁴See Footnote 2.

It may be useful to describe some special game theoretic features of distributed games. Such games can be described as finitely repeated games with incomplete information about the order of stages with the additional feature that the level of information monitoring is strategically controlled by the players via the messages. Moreover, a distributed game has imperfect recall; The use of agents enable the players to commit themselves not to remember facts that they have already known⁵.

In the next section we apply our model to the problem of cooperation in the distributed Prisoner's Dilemma game. Our definition of distributed games can be modified in several natural ways in order to deal with other types of distributed interactions.

3. Cooperation in the Distributed Prisoner's Dilemma Game. In this section we analyze a distributed game with two players and n locations, in which at each location $i = 1, 2, \dots, n$, the agents play the Prisoner's Dilemma game described below:

$$\begin{array}{cc} & D & C \\ D & (a, a) & (b, 0) \\ C & (0, b) & (c, c) \end{array},$$

where $b > c > a > 0$. The probability distribution λ assigns probability $\frac{1}{n!}$ to each order of plays. It is assumed that the message space contains only one real message m_1 . That is $M^1 = M^2 = \{m_0, m_1\}$, where m_0 means "no message". We show that under Condition (3.3) (below), there exists an equilibrium $\sigma = (\sigma_1, \sigma_2)$ which induces the outcome (C, C) at each location. Indeed, for $j = 1, 2$, define $\sigma^j = ((f_i^j, g_i^j))_{i=1}^n$ as follows: f_i^j assigns the action C to the history of no previous $n-1$ times messages $h = \overbrace{[(m_0, m_0), \dots, (m_0, m_0)]}$, and it assigns the action D to any other history of messages. In addition, independently of the history of messages $h \in H_{-i}$, $g_i^1(h, (C, C)) = g_i^1(h, (D, C)) = m_0$, $g_i^1(h, (C, D)) = g_i^1(h, (D, D)) = m_1$, and $g_i^2(h, (C, C)) = g_i^2(h, (C, D)) = m_0$, $g_i^2(h, (D, C)) = g_i^2(h, (D, D)) = m_1$. Obviously if all agents obey the strategy profile σ , then the outcome (C, C) is played in every location. We proceed to show that for $j = 1, 2$, Player j does not have a profitable deviation (if Condition (3.3) is satisfied). As the players are symmetric we discuss only deviations of Player 1. Assume Player 1 deviates at k locations, $1 \leq k \leq n$. Without loss of generality we assume that all these deviations involve the action selecting strategies and not the message selecting strategies. Moreover, it can be

⁵ Recently, following Piccione and Rubinstein (1995), games with imperfect recall have been extensively discussed in the literature. See e.g., Aumann, Hart and Perry (1996, 1997) and Halpern (1996).

assumed that these deviations have the following form: the agents of Player 1 use D at k of the locations, independently of the history of messages. The other agents of Player 1 use their part of the strategy σ^1 . If the first defection occurs at stage t , $1 \leq t \leq n$, then Player 1 receives $t - 1$ times the value c , one time the value b and $n - t$ times the value a . Therefore the expected payoff of Player 1, given that the first defection occurs at stage t is:

$$(3.1) \quad v_t = (t - 1)c + b + (n - t)a.$$

Hence, the expected payoff upon a deviation is

$$(3.2) \quad v^k = \sum_{t=1}^n v_t p_t^k,$$

where p_t^k is the probability that the first defection occurs at stage t , given that it defects at k locations.

We now show that the expected value of Player 1 resulting from a deviation in k locations is non-increasing in k . That is, $v^k \geq v^{k+1}$ for $1 \leq k \leq n - 1$. Indeed, denote by F_t^k the probability of first defection at stage t or before. That is,

$$F_t^k = \sum_{s=1}^t p_s^k.$$

It is easily verified that F^{k+1} stochastically dominates F^k (that is, $F_t^k \leq F_t^{k+1}$ for every $1 \leq t \leq n$), and that $(v_t)_{t=1}^n$ is a non-decreasing sequence. Therefore (3.2) implies the desired monotonicity. Hence, it suffices to show that a deviation at precisely one location is not profitable, that is, $v^1 \leq nc$. Since $p_t^1 = \frac{1}{n}$ for all t , we get from (3.1) and (3.2) that

$$v^1 = \frac{1}{n} \sum_{t=1}^n v_t = \frac{1}{n} \sum_{t=1}^n (t(c - a) + na - c + b).$$

It follows that the equilibrium condition is

$$\frac{n + 1}{2}(c - a) + na - c + b \leq nc.$$

That is,

$$(3.3) \quad b \leq c + \frac{n - 1}{2}(c - a).$$

Thus, the cooperation outcome at each location can be obtained in equilibrium if the number of locations is sufficiently large. The precise meaning of ‘‘sufficiently large’’ is given by (3.3).

4. Cooperation in the Prisoner's Dilemma Game- Related Literature. It is instructive to compare our result with other cooperation results for the Prisoner's Dilemma game, obtained by assuming finite sequentiality and bounded rationality⁶. Radner (1986) showed that when the players are ready to use ε -equilibrium, then when the number of repetitions increases the corresponding sets of ε -equilibria allow longer and longer period of cooperation. Kreps, Milgrum, Roberts, and Wilson (1982) showed that if we assume that with an arbitrary small but positive exogenous probability, one of the players is playing tit-for tat rather than maximizing, then with sufficiently long repetition, all sequential equilibria outcomes are close to cooperative (see, however, Fudenberg and Maskin (1986)). Aumann and Sorin (1989) proved that when every player ascribes a small positive exogenous probability to his opponent being an automaton with bounded recall, then every equilibrium in sufficiently long repetition of the game is close in the payoff space to the cooperative outcome. Neyman (1985) deals with finitely repeated Prisoner's Dilemma game in which the players are restricted to use automata with a fixed number of states. When the number of stages is large relative to the number of states⁷, the automata cannot effectively count the number of previous stages, and as in our case, this ignorance regarding the number of stages enables cooperation (in all stages) in equilibrium. Hence in the automaton model the players are not able to process the available information, while in our model, the players jointly and strategically decide not to obtain this information. We think that our model is more realistic in the distributed systems (e.g., Internet) setup in which it is assumed that players broadcast messages and act in a parallel asynchronous setting, while controlling software agents with a tremendous counting ability. Zemel (1989) use the finite automata model of Neyman with the additional feature of allowing the players to send and receive messages. This additional feature allows cooperation by saturating the computational resources of the players and thus preventing them from utilizing complex strategies.

The recent paper of Neyman (1996) is related to our work. It is proved there that the players in a repeated Prisoner Dilemma can (almost) reach the cooperation outcome if they do not share common knowledge regarding the true number of stages. Roughly speaking, our model exhibits a "real life" situation where such

⁶The folk theorems (See e.g., Aumann and Shapley (1994), Rubinstein (1979)) show that cooperation is possible in equilibria of the *infinite* repeated game.

⁷Or relative to the number of states to the power of $\frac{1}{\varepsilon}$, when we are ready to settle for ε -equilibrium.

a lack of common knowledge is possible. Note however that this lack of common knowledge is not exogenous, but it is the result of the strategic decisions of the players to let their agents and opponents be ignorant about the number of (future) stages.

REFERENCES.

Aumann, R.J. and Sorin, S. (1989), Cooperation and Bounded Recall, *Games and Economic Behavior*, 1, 5–39.

Aumann, R. J. and Shapley, L.S. (1994), Long Term Competition– A game Theoretic Approach, in *Essays in game Theory, In Honor of Michael Maschler*, edited by N . Meggido, 1–15.

Aumann, R. J., Hart, S. and Perry, M. (1996), The Absent-Minded Driver, *Center for Rationality and Interactive Decision Theory*, DP No. 94R.

Aumann, R. J., Hart, S. and Perry, M. (1997), The Forgetful Passenger, *Center for Rationality and Interactive Decision Theory*, DP No. 130.

CACM (1994), Intelligent Agents, *Special Issue of CACM (Communications of the ACM)*, 37(7).

Diffie, W. and Hellman, M.E. (1976), New Directions in Cryptography, *IEEE Transactions on Information Theory*, 22(5), 644–654.

Fudenberg, D. and Maskin, E. (1986), The Folk Theorem in Repeated games with Discounting and Incomplete Information, *Econometrica*, 54, 533–554.

Halpern, J. Y. (1996), On Ambiguities in the Interpretation of Games Tree *Proceeding of the Sixth TARK Conference*, 77–96.

Kreps, D., Milgrum, P., Roberts, J. and Wilson, R. (1982). Rational Cooperation in Finitely-Repeated Prisoners' Dilemma, *Journal of Economic Theory*, 27, 245–252.

Neyman, A. (1985), Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoner's Dilemma, *Economics Letters*, 19, 227–229.

Neyman, A. (1996), Cooperation in Repeated Games when the Number of Stages is not Commonly Known, *Center of Rationality Working Paper*.

Piccione, M. and Rubinstein, A. (1995), On the Interpretation of Decision Problems with Imperfect Recall, *The Sackler Institute for Economic Studies, Tel Aviv*

University, Working Paper No. 24-94.

Radner, R. (1986), Can Bounded Rationality Resolve the Prisoners' Dilemma, in *essays in Honor of Gerard Debreu* Edited by A. Mas-Colell and W. Hildenbrand, North-Holland Publ., Amsterdam. Epsilon-Equilibria

Rubinstein, A. (1979), Equilibrium in Supergames With the Overtaking Criterion, *Journal of Economic Theory*, 21, 1–9.

Tanenbaum, A.S. (1988), *Computer Networks, 2nd Edition*, Prentice Hall.