Available online at www.sciencedirect.com
SCIENCE DIRECT
GAMES and Economic Behavior
Games and Economic Behavior 45 (2003) 114–131
www.elsevier.com/locate/geb

A computer scientist looks at game theory

Joseph Y. Halpern [1]

Cornell University, Computer Science Department, Ithaca, NY 14853, USA

Received 15 February 2001

Abstract

I consider issues in distributed computation that should be of relevance to game theory. In particular, I focus on (a) representing knowledge and uncertainty, (b) dealing with failures, and (c) specification of mechanisms.

JEL classification: D80; D83

1. Introduction

There are many areas of overlap between computer science and game theory. The influence of computer science has been felt perhaps most strongly through complexity theory. Complexity theory has been viewed as a tool to help capture bounded rationality, going back to the work of Neyman (1985) and Rubinstein (1986). In addition, it is well understood that complexity-theoretic notions like NP-completeness help categorize the intrinsic difficulty of a problem. Thus, for example, a result showing that, even in simple settings, the problem of optimizing social welfare is NP-hard (Kfir-Dahav et al., 2000) shows that the standard procedure of applying Clarke's mechanism, despite its many benefits, is not going to work in large systems.

Perhaps less obvious is the interplay between game theory and work in distributed computing. At the surface, both areas are interested in much the same problems: dealing with systems where there are many agents, facing uncertainty, and having possibly different

E-mail address: halpern@cs.cornell.edu.
URL: http://www.cs.cornell.edu/home/halpern.

[1] This paper is based on an invited talk I gave at Games 2000 in Bilbao, Spain. Supported in part by NSF under grants IRI-96-25901 and IIS-0090145 and by ONR under grants N00014-00-1-03-41, N00014-01-10-511, and N00014-01-1-0795.

0899-8256/$ – see front matter
doi:10.1016/S0899-8256(02)00529-8

goals. In practice, however, there has been significant difference in emphasis in the two areas. In distributed computing, the focus has been on problems such as fault tolerance, scalability, and proving correctness of algorithms; in game theory, the focus has been on strategic concerns (that is, playing so as to optimize returns, in light of the preferences of other agents). In this paper, I hope to make the case that each area has much to learn from the other. I focus on three particular topics:

- the representation of games (and, in particular, the knowledge and uncertainty of players in a game),
- strategic concerns vs. fault tolerance, and
- specification of mechanisms.

The following sections deal with each of these topics in turn.

2. Representing games as systems

In order to analyze a game, we must first represent it. The two most common representations in the literature are the normal-form representation and the extensive-form representation. As is well known, the extensive-form representation brings out the temporal aspects of the game better, as well as explicitly representing (at least some aspects) of the players' knowledge. Consider the game that is represented in Fig. 1 in both normal form and extensive form. The extensive-form representation brings out clearly that the game takes place over time, with the first player's second move, for example, occurring after the second player's first move. Moreover, when the first player makes that second move, he does not know what the second player's move is. However, as is also well known, the information sets used in the extensive-form representation do not capture all aspects of a player's information. For example, they cannot be used to capture beliefs one player has about what strategy the other player is using, or notions like rationality and common knowledge of rationality. The state-space representation does better in this regard.

2.1. The state-space representation

The state-space representation, first used in the economics literature by Aumann (1976), is actually a variant of the standard possible-worlds model for knowledge in the philosophical literature that goes back to Hintikka (1962); see (Fagin et al., 1995,

	A	D
aa	(3,3)	(4,2)
ad	(3,2)	(4,2)
da	(1,2)	(1,3)
dd	(1,2)	(1,3)

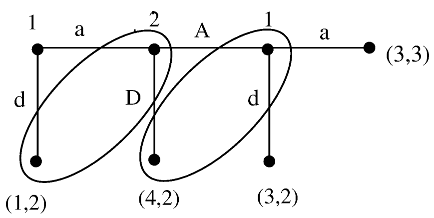

Fig. 1. Representing a game in normal form and in extensive form.