Distributed Computing Meets Game Theory

Joseph Y. Halpern
Cornell University, USA

Abstract

The question of whether a problem in a multiagent system that can be solved with a trusted mediator can be solved by just the agents in the system, without the mediator, has attracted a great deal of attention in both computer science (particularly in the cryptography community) and game theory. In cryptography, the focus on the problem has been on secure multiparty computation, where each agent has some private information and the agents want to compute some function of this information without revealing it. This can be done trivially with a trusted mediator: the agents just send their private information to the mediator, who computes the function value and sends it to all of them. Work on multiparty computation conditions under which this can be done without a mediator, under the assumption that at most a certain fraction of the agents are faulty, and do not follow the recommended protocol. By way of contrast, game theory is interested in implementing mediators using what is called "cheap talk", under the assumption that agents are rational. We are interested in combining both strands: We consider games that have (k,t)-robust equilibria when played with a mediator, where an equilibrium is (k,t)-robust if it tolerates deviations by coalitions of rational players of size up to k and deviations by up to t players who can be viewed as faulty (although they can equally well be viewed as rational players with unanticipated utilities). We prove matching upper and lower bounds on the ability to implement such mediators using cheap talk (that is, just allowing communication among the players). The bounds depend on (a) the relationship between k, t and n, the total number of players in the system; (b) whether players know the exact utilities of other players; (c) whether there are broadcast channels or just point-to-point channels; (d) whether cryptography is available; and (e) whether the game has a (k+t)-punishment strategy; that is, a strategy that, if used by all but at most k+t players, guarantees that every player gets a worse outcome than they do with the equilibrium strategy.