Game-Theoretic Aspects of Distributed Computing

Chryssis Georgiou
Dept. of Computer Science, University of Cyprus
1678 Nicosia, Cyprus
chryssis@cs.ucy.ac.cy

The Workshop on Game-Theoretic Aspects of Distributed Computing took place on September 26th and it was co-located with DISC 2009[1]. The workshop was organized by Chryssis Georgiou (University of

[1] http://disc2009.gsyc.es/workshops
[1] http://disc2009.gsyc.es

Figure 1: Some of the workshop participants. Picture courtesy of Antonio Fernandez Anta.

Cyprus, Cyprus) and Paul Spirakis (RACTI and University of Patras, Greece) and it was sponsored by the University of Cyprus[2] and the Comunidad de Madrid[3]. Further details can be obtained from the Workshop's website[4].

Workshop's Objective

In traditional Distributed Computing, the behavior of the system components (i.e., processors/processes/nodes/agents) is characterized *a priori* as either "good" or "bad", depending on whether they follow the prescribed protocol or not. In Game Theory, players are assumed to be *rational* or *selfish*, that is, they act on their own self-interest and they do not have an *a priori* established behavior. In other words, the players decide on how to act in an attempt to increase their own benefit (a quantified measure).

Game Theory has long been considered in many fields, ranging from Economics to Law Enforcement and Voting Decision. With the evolution of the Internet, Game Theory has found many applications in Networks, and in Distributed Computing in general. Such examples include Internet Routing, Resource/Facility Location and Sharing, Containment of Viruses Spreading, Secret Sharing, and Web-based Task Computations.

The purpose of the workshop was to present recent works that consider Distributed Computing issues from a Game-Theoretic view and approach. The *main objective* of the workshop was to enable members of the Distributed Computing community (especially students and junior researchers) to realize the potential of what it can be called *Game-Theoretic Distributed Computing*.

The workshop included five invited talks delivered by prominent researchers working on Game-Theoretic Distributed Computing issues: *Ioannis Caragiannis* (University of Patras, Greece), *Stefano Leonardi* (Sapienza University of Rome, Greece), *Giuseppe (Pino) Persiano* (Universita di Salerno, Italy), *Maria Serna* (Universitat Politecnica de Catalunya, Spain), and *Paul Spirakis* (RACTI and University of Patras, Greece). Figure 1 shows the workshop's invited speakers, organizers and some of the participants.

Overview of the Invited Talks

We give an overview of the five invited talks in the order they were presented.

[2] http://www.ucy.ac.cy
[3] http://www.madrid.org
[4] http://sites.google.com/site/wgtadc

***Strategic Games to Analyze Unreliable Systems*, by Maria Serna.** Being the first speaker, Maria first presented some fundamental terminology of Game Theory (including strategic games, pure and mixed strategies, utilities and cost functions, and Nash equilibria). Then she proceeded to the main topic of the talk, which is a work performed jointly with J. Gabarro, A. Garcia, P. Kilpatricka and A. Stewart. They formulated and analyzed, from a game-theoretic view, two problems arising in web or grid environments with unreliable behavior: service failures during grid orchestrations and resource allocation in grids.

Maria first presented the idea behind the problems' formulation. A service failure may be catastrophic in that it causes an entire grid application to fail. Alternatively, a grid manager may utilize alternative services in the case of a failure, allowing an orchestration to recover. The proposed approach attempts to provide an alternative for the case of bounded number of failures in between optimistic and pessimistic situations and encapsulates the idea of *limited failures* and *destructive* versus *not destructive* behavior. The talk concentrated in dealing with failures prior to the execution of the tasks; such failures can be anticipated and modeled by what they call *risk profile*. A risk profile is the mean of modeling situations in a way that is neither overly optimistic nor overly pessimistic. Risk profiles provide the opportunity of defining a class of games that they call *Angel and Daemon* games, in which the angel controls non-malicious failures while the daemon controls malicious failures. Then Maria presented the two problems under investigation and showed how they can be modeled as Angel-Daemon games. She concluded the talk with the presentation of some results on the existence of pure Nash equilibria and the complexity of the corresponding computational problems.

***Efficient Coordination Mechanisms for Selfish Scheduling*, by Ioannis Caragiannis.** Ioannis presented *coordination mechanisms* for scheduling selfish jobs on unrelated machines. As Ioannis explained, a coordination mechanism aims to mitigate the impact of selfishness of jobs on the efficiency of schedules by defining a local scheduling policy on each machine. These scheduling policies induce a game among the jobs and each job prefers to be scheduled on a machine so that its completion time is minimum given the assignments of the other jobs; the maximum completion time among all jobs is considered as the measure of the efficiency of schedules. The approximation ratio of a coordination mechanism quantifies the efficiency of pure Nash equilibria (price of anarchy) of the induced game.

The talk covered simple strongly local and non-preemptive coordination mechanisms (such as ShortestFirst, LongerstFirst, and Randomized) as well as recent preemptive ones (such as ACOORD, BCOORD, and CCOORD) that achieve logarithmic or polylogarithmic (in terms of the number of the machines) approximation ratio. These preemptive coordination mechanisms make use of *m-efficient assignments*. In such assignments, a job is never assigned to a machine where its inefficiency is more than m. The talk concluded with a discussion on open questions. For example, is it possible to obtain constant approximation ratio? Is the knowledge of m really necessary? How about mixed Nash equilibria?

***Routing Traffic Through Selfish Capacitated Links*, by Giuseppe (Pino) Persiano.** Pino's talk considered the problem of *routing traffic* over communication links that are owned by selfish agents; agents are selfish in the sense that might misreport link characteristics to the routing algorithm so to induce a better (for the agent) allocation of traffic to links. As Pino pointed out, this is a classical problem in *mechanism design* where one seeks payments for the agents so to make it profitable to report the true characteristics of the links (that is, make the agents *truthful*), and thus the routing algorithm can work on real data.

In the talk, Pino presented a recent work, performed jointly with V. Auletta and P. Penna, where they study the novel scenario in which links have traffic capacities. As Pino explained in detail, the obtained results are in stark contrast with the known results for another important case of agents with capacities (bidders with budget constraints). Several reseach questions on this topic remain open, for some of which Pino gave insight on possible directions in addressing them.

***Utilitarian Mechanism Design for Multi-Objective Optimization*, by Stefano Leonardi.** Stefano presented recent work, performed jointly with F. Grandoni, P. Krysta and C. Ventre, on algorithmic mechanism design for *multi-objective optimization* problems. In a classic optimization problem the complete input data is assumed to be known to the algorithm. As Stefano pointed out, this assumption may not be true anymore in optimization problems motivated by the Internet where part of the input data is private knowledge of independent selfish agents. The goal of *algorithmic mechanism design* is to provide (in polynomial time) a solution to the optimization problem and a set of incentives for the agents such that disclosing the input data is a dominant strategy for the agents. In case of NP-hard problems, the solution computed should also be a good approximation of the optimum.

In the talk, Stefano focused on mechanism design for multi-objective optimization problems. In this setting a main objective function is given, along with a set of secondary objectives which are modeled via budget constraints. Multi-objective optimization is a natural setting for mechanism design as many economical choices ask for a compromise between different, partially conflicting, goals. The main contribution of this work (which Stefano explained rigorously) is showing that two of the main tools for the design of approximation algorithms for multiobjective optimization problems, namely *approximate Pareto curves* and *Lagrangian relaxation*, can lead to truthful approximation schemes.

***Learning and Enforcement in Competitive Environments*, by Paul Spirakis.** Paul's talk considered the problem of efficiently *learning and enforcing* equilibria strategies in competitive environments. Paul first presented basic notions of Machine Learning and discussed what can be learned in competitive environments consisting of selfish agents. In particular, he showed that *correlated equilibria strategies* can be learned efficiently in such environments. Furthermore, he demonstrated that Nash equilibria can be learned in the special case of zero sum bimatrix games. Then the talk focused on *repeated games* of many agents and discussed the issue of collaborative enforcement of individual behaviours that lead to equilibria play. The notions of *regret*, of a *threat point* and of a *correlated threat point* were defined and explained through several examples. The talk concluded by demonstrating how the notion of a correlated threat point can help in efficient enforcement of equilibria strategies.

Acknowledgements

Paul Spirakis and I would like to thank the DISC'09 Organization Committee chairs, *Antonio Fernández Anta* (Universidad Rey Juan Carlos, Spain) and *Vicent Cholvi Juan* (Universitat Jaume I, Spain) for their cordial help and support. We especially appreciate their involvement in obtaining sponsorship from the Comunidad de Madrid. Also, I would like to thank *Idit Keidar* (Technion, Israel) for inviting me to write this report.