Brief Announcement: Game Authority for Robust and Scalable Distributed Selfish-Computer Systems*

Shlomi Dolev,† Elad M. Schiller,‡ Paul G. Spirakis,§ and Philippas Tsigas‡

† Dep. Computer Science
Ben-Gurion University of the Negev, Israel
dolev@cs.bgu.ac.il

‡ Dep. Computer Science
Chalmers University of Technology, Sweden
{elad,tsigas}@chalmers.se

§ Research Academic
Computer Technology Institute
Greece
spirakis@cti.gr

Categories & Subject Descriptors: C.2.4 [Distributed Systems]: Distributed applications
General Terms: Algorithms, Design, Reliability, Theory.
Keywords: Game Theory, Byzantine Agreement.

Game theory analyzes social structures of agents that have freedom of choice within a moral code. The society allows freedom and selfishness within the moral code, which social structures enforce, i.e., legislative, executive, and judicial. Social rules encourage individual profit from which the entire society gains. Distributed computer systems can improve their scalability and robustness by using explicit social structures. We propose using a game authority middleware for enforcing the moral code on selfish agents.

The power of game theory is in predicting the game outcome for specific assumptions. The prediction holds as long as the players cannot tamper with the social structure, or change the rules of the game, i.e., the prisoner cannot escape from prison in the classical prisoner dilemma. Therefore, we cannot predict the game outcome without suitable assumptions on failures and honest selfishness.

There are attempts to define various aspects of selfish-computer systems: selfish MAC Layer that does not back off in [3], Byzantine Nash Equilibrium of replicated state-machine in [1], and selfish mechanism for virus inoculation in the presence of malicious agents in [8], to name a few. In fact, [8] discovers that the performance ratio between selfish mechanisms that do and do not have malicious agents (named the *price of malice*, PoM), are as important as a performance ratio between selfish mechanisms and centralistic mechanisms, such as the *price of anarchy* (PoA) [7] and the *price of stability* (PoS) [2]. We study these performance ratios in the context of explicit game implementation and discover significant improvements.

We argue that when designing distributed selfish-computer systems it is unsuitable to assume that all software layers and components act selfishly. Under this strong assumption, the designer has to consider a complex game among all selfish agents, which have many possible software actions and imprecise utility in the presence of failures. Moreover, not all games have a predictable outcome; many games have very long stabilization periods and incomplete information games[1] deteriorate system efficiency. Consequently, designers cannot predict the outcome without a suitable perspective on the various system aspects.

The middleware services. The game authority facilitates interaction among agents of the application layer, where users control programs. The middleware implements a social structure using none-selfish components (some of which can be Byzantine.)[2] Our social structure follows the principle of power separation. The key services are:

• **The legislative service** which allows agents to set up the rules of the game in a democratic manner, e.g., robust voting, which results in the majority preferable outcome.

• **The judicial service** which audits the agents' actions and orders the executive service to punish agents following their foul play.

• **The executive service** which executes actions and manages their associated information: publish utilities, collect choice of actions, and play outcome. Moreover, by order of the judicial service, this service restricts the action of dishonest agents.

The middleware properties. This work shows the existence of the game authority and its middleware services that can achieve the following properties:

• *Scalability with a freedom of choice.* The legislative service allows a society of self-consciousness agents to vote for their preferable game and avoid playing games that compromise performances beyond the necessary extent for providing the freedom of choice. We introduce a performance ratio between the novel democratic social structures and the classic dictatorship designs in which there is a predefined program executed by all non-Byzantine processes. We name this new criterion the *multi-round anarchy cost*. We say that an agent is honestly selfish if it rationally follows the legislated law. We present an asymptotic similarity between centralistic solutions and a well-chosen game among honestly selfish agents for resource sharing.

*The first author is partially supported by Israeli Ministry of Trade & Industry, & the Rita Altura Trust Chair in Computer Sciences. The third author is partially supported by the EU 6th Framework Programme under contract 001907 (DELIS). A preliminary version (together with the correctness proofs) appears in [5].

PODC'07, August 12–15, 2007, Portland, Oregon, USA.
ACM 978-1-59593-616-5/07/0008.

[1] Game theory uses this term to describe a game where individual agents may not be able to predict (precisely) the effect their actions will have on the other agents.

[2] Assuming standard requirements for Byzantine agreement, i.e., more than $2/3$ of the processes are (selfish but) honest, and authentication utilizes a Byzantine agreement that needs only a majority. Moreover, the communication graph is not partitioned, e.g., there are $2f + 1$ vertex disjoint paths between any 2 processes, in the presence of at most f Byzantine processes.

• *Robust assurance of honest selfishness.* The judicial service emulates games, monitors agents, and blocks the dishonest ones. We use cryptographic techniques for auditing agents and preserving their privacy.

An essential design criterion is Byzantine fault tolerance. We show that for complete information games, the judicial service can base its implementation on Byzantine agreement, assuming the existence of a trusted executive service. Moreover, we consider a self-stabilizing implementation of the game authority that is based on a new self-stabilizing Byzantine agreement algorithm.

Our contribution. We design the first game authority middleware that enforces the rules of a complete information game. Our design is inspired by the power separation principle by which the legislative service votes for the rules of the game, the judicial service audits the agents' selfish choices, and the executive service carries out the agents' actions. The judicial service detects agents that do not select their action according to the rules made by the legislative service and then the executive service disqualifies them from playing. We base our design on a self-stabilizing Byzantine agreement according to which the majority of participants audit the game while dealing with unexpected faults. We use several other cryptographic techniques for auditing mixed strategies while preserving the privacy of the agents.

We show that the game authority can guarantee asymptotically optimal resource allocation performances according to our new and suitable cost criterion. We further estimate the system performances using existing criteria and show how to replace the PoA with the PoS.

Analysis. The literature compares the performance of systems that are restricted to rough consensus with the best centralistic solution. e.g., the price of anarchy [7] considers the worst Nash Equilibrium and the price of stability [2] considers the best Nash Equilibrium. There is a need for a new cost criterion that better captures the system where selfish agents legislate the rules of the game, enforce rational actions, and repeat the play. We say that an agent is honestly selfish if it rationally follows the legislated law. We call the new criterion *multi-round anarchy cost* because it estimates the performances of honestly selfish agents that continuously play their preferable game.

We now turn to exemplifying the intensity of the game authority, which allows agents to choose their preferred game. The *repeated resource allocation* (RRA) is a scheduling problem with demand bursts, limited coordination, and transient fault disruption. In every play, each agent places a (single unit) demand for a resource. We assume that at the end of every play all agents know the load that exists on the resources, that the number of plays is *unknown* (i.e., every play could be the last one). Thus, selfish agents choose resources according to the Nash equilibrium.

THEOREM 1. *The game authority that supervises the RRA game guarantees multi-round anarchy cost in* $O(1)$.

Consider a concrete RRA implementation in which a consortium of Internet companies share licenses for advertisement clips on video Web sites. We note that the unpredictable loads on the hosts cause service availability issues. There are many complex ways to model this scenario. Theorem 1 shows that the simplest way is optimal. Therefore, it could be that the consortium majority prefers backlog size as the host's only selection criterion (and reject criteria such as video content and attempts for synchronized advertisement).

Reduction to Byzantine agreement protocol. Complete information games have public agents' knowledge and state. In [5] we present an algorithmic scheme for implementing the game authority for complete information games. The algorithm relies on a Byzantine agreement protocol (BAP) and two cryptographic primitives: a commitment scheme and pseudo-random generator.

The game authority uses a Byzantine common pulse (e.g., [4]) to repeat the sequence of: (1) verifying that the majority of agents are interested in playing the game (by allowing them to repeatedly agree on a game) using the legislative service, (2) conducting an audition for the next play (in order to find all the foul plays) using the judicial service, and (3) executing the play (and learning the outcome) using the executive service. In the full version of the paper [5], we show that the algorithm can safely emulate any complete information game.

THEOREM 2. *The game authority can secure the execution of complete information games.*

Selfish-stabilizing. We present a self-stabilizing Byzantine agreement (SSBA) algorithm. We name the combined system design selfish-stabilizing, because the system can recover from transient failures and short-lived myopic logic, say, due to transient computational resource shortage.

The algorithm is a composition of two distributed algorithms. We use the self-stabilizing Byzantine clock synchronization algorithm of [4]. Whenever the clock value reaches the value 1, the SSBA algorithm invokes the BAP of (e.g., [6]). We take the clock size, $\log M$, to be large enough to allow exactly one Byzantine agreement.

THEOREM 3. *SSBA is a self-stabilizing BAP.*

Conclusions. The game authority is a general-purpose middleware that can protect honestly selfish agents when playing a complete information game. The agents are running in the application layer, chose their starting PSP after electing their preferable game to play. One of the key benefits of the game authority is that it replaces the cost of anarchy with the cost of stability. Moreover, we show that the middleware has a self-stabilizing implementation.

References

[1] A. S. Aiyer, L. Alvisi, A. Clement, M. Dahlin, J.-P. Martin, and C. Porth. Bar fault tolerance for cooperative services. *SIGOPS Oper. Syst. Rev.*, 39(5):45–58, 2005.

[2] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *FOCS '04: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS'04)*, pp. 295–304, Washington, DC, USA, 2004. IEEE Computer Society.

[3] M. Cagalj, S. Ganeriwal, I. Aad, and J.-P. Hubaux. On selfish behavior in csma/ca networks. In *INFOCOM*, pp. 2513–2524. IEEE, 2005.

[4] S. Dolev and J. L. Welch. Self-stabilizing clock synchronization in the presence of byzantine faults. *J. ACM*, 51(5):780–799, 2004.

[5] Dolev Shlomi, Schiller Elad, Spirakis Paul. Game authority: for robust distributed selfish-computer systems. Technical report, DELIS, 2006. Accessible via http://delis.upb.de/docs/.

[6] J. A. Garay and Y. Moses. Fully polynomial byzantine agreement for processors in rounds. *SIAM J. Comput.*, 27(1):247–290, 1998.

[7] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In C. Meinel and S. Tison, editors, *STACS*, volume 1563 of *Lecture Notes in Computer Science*, pp. 404–413. Springer, 1999.

[8] T. Moscibroda, S. Schmid, and R. Wattenhofer. When selfish meets evil: byzantine players in a virus inoculation game. In *PODC '06: Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, pp. 35–44, New York, NY, USA, 2006. ACM Press.