Emergent Socially Rational Utilitarianism in a Peer-to-Peer System[1]

Andrea Marcozzi, David Hales {marcozzi, hales}@cs.unibo.it
University of Bologna, Italy
Poster Abstract Submission for ECCS2006

For many applications peer-to-peer (P2P) systems require their member nodes (or agents) to behave in a socially beneficial (non-egotistical) way. Kalenka and Jennings (1999) termed this requirement as the "Principle of Social Rationality": if an agent has a choice of actions it should chose the action that maximizes the social utility (sum of all agent utilities in the system). This principle can be contrasted with classical individual rationality that states agents should select actions that maximize their individual utility.

However, developing protocols for realistic P2P systems that adhere to the principle of social rationality is very difficult and potentially so costly as to negate the benefits. This is because P2P systems have no central control, are potentially huge (composed of millions of nodes) and have high node turnover (with users continually entering and leaving the system). In addition, selfish or malicious nodes can get into the system via hacked client programs. These factors mean that individual nodes, even if they wish to follow a socially rational principle, often will not have enough information to gauge the effects of their actions on others. It would be too costly for every node to ask every other node to report its state before every individual action – even if it was possible.

Recently, simple locally adaptive protocols have been proposed (Marcozzi 2005, Hales 2005) that claim to produce socially rational outcomes through a process of self-organization even though nodes only act on their own utility values. In this approach nodes preferentially copy other nodes (by duplicating their behaviour and links) that have higher utilities. However, in these previous works only specific scenarios are considered in which certain plausible utility values are selected. This method can be contrasted with those relating to the field of mechanism design giving concrete ways to reduce the lack of coordination (Koutsoupias 1999). In this poster we take one such existing P2P model (the SkillWorld – Hales 2005) and modify the utilities to explore a large space of possible values. In each case we checked if the protocol maximized the collective utility or not. We found in each case that if the collective cost of an action was less than or equal to the collective benefit the protocol self-organized the network to a state where nodes selected this action. We present a synthesis of results from computer simulations covering over ten billion interactions.

References

A. Marcozzi, D. Hales, G. Jesi, S. Arteconi, O. Babaoglu. *Tag-Based Cooperation in Peer-to-Peer Networks with Newscast*. In Self-Organization and Autonomic Informatics (I), edited by: H. Czap, R. et al. IOS Press, Netherlands. 2005.

D. Hales. *Emergent Group-Level Selection in a Peer-to-Peer Network*. Proceedings of the European Conference on Complex Systems (ECCS'05), Nov. 14th, Paris, France. Publishers: i6doc, Belgium. 2005.

D. Hales and S. Arteconi. *SLACER: A self-organizing protocol for coordination in peer-to-peer networks*. IEEE Intelligent Systems, 21(2):29–35, Mar/Apr 2006.

S. Kalenka and N. R. Jennings. *Socially Responsible Decision Making by Autonomous Agents. Cognition, Agency and Rationality*. (eds. Korta, K., Sosa, E., Arrazola, X.) Kluwer 135-149, 1999.

E. Koutsoupias and C. H. Papadimitriou. *Worst-case Equilibria*. STACS'99.

[1] Partially supported by the EU within the 6th Framework Programme under contract 001907 Dynamically Evolving, Large Scale Information Systems (DELIS).