

Sociologically Inspired Approaches for Self-^{*}: Examples and Prospects

David Hales

The University of Bologna, Italy.
dave@davidhales.com,
WWW home page: <http://davidhales.com>

Abstract. One way of approaching the engineering of systems with self-^{*} properties is to examine naturally occurring systems that appear to have such properties. One line of work examines biological theories and phenomena. Ideas from the social sciences are less well explored as a possible source of self-^{*} techniques. We briefly overview some recent work that follows this latter approach and consider some specific prospects for future work.

1 Why Social Science?

Human social systems appear to be scalable, self-repairing and self-regulating and often robust. They spontaneously form, and emerge apparently functional structures, institutions and organisations.

Much social scientific research has been produced concerning why and how social phenomena occur and social science itself has numerous sub-disciplines, sub-schools, methodologies and approaches.

We believe that many of the deep engineering problems inherent in the self-^{*} approach can be thought of as sociological questions.

Recently, new computational approaches have been applied to explore the complex processes of emergence that often characterise social phenomena. This approach forces a new kind of rigour on social theory construction and offers the prospective self-^{*} engineer a possible source of ideas to plunder.

2 Computational Social Science

It is only very recently, with the arrival of cheap, fast, desktop computers and social science researchers who know how to program them, that a new area of ‘computational social science’ has begun to emerge.

There has been an explosion of published work concerning sociologically motivated computational models [5–7, 14]. In contrast to early equation-based ‘high-level’ models, in which there was no space of individual behaviours, much of these models are described as ‘agent-based’.

Agent-based modelling in these contexts means a discreet, individual and event-based approach. Individual behaviours of agents (representing people, groups

or institutions) are programmed explicitly as a computer program. A population of such agents (or programs) inhabiting a shared environment are then allowed to interact over time and the emergent results and outcomes are observed. It is therefore a prerequisite of such work that agent behaviours must be specified algorithmically.

The emphasis of much computational social science is on the emergent properties of these ‘artificial societies’. By experimentation and observation researchers attempt to gain general insights into mechanisms of social emergence and then to relate these to real human societies.

It should be noted that the relationship between real social systems and computer models is, and probably always will be, highly controversial — human social systems are so complex, fluid and political (by definition) that debates about what constitutes adequate validation and verification of models rarely converge to agreement. However, these kinds of debates do not need to trouble an engineer looking for new techniques to construct self-* systems.

3 A Brief Note on Game Theory

Some branches of economics, particularly classical game theoretical approaches, formalised their subject matter, analytically, some time ago. This was due, in part, to the advances made by von Neumann and Morgenstern’s seminal work [23] and early pioneers such as Nash [18].

However, due to the focus and strong assumptions of classical game theory — quite proper for the original focus and application of the work — a lot of results are hard to apply to typical self-* scenarios (e.g. noisy, dynamic and with little information concerning the possible behaviour of other units in the system). The classical approach gives analytical proofs of the ‘best’ way to act in a given situation under the assumption that each actor or agent has complete information and infinite computational resources.

Despite these qualifications, classical game theoretical analysis has many possible areas of application [3] — but we will not concentrate on these here. Also the abstracted scenarios (games) constructed by game theorists to capture certain kinds of social interactions are useful as a basis for evaluating other kinds of modelling techniques (as we shall see later with the Prisoner’s Dilemma game).

Interestingly, within economics there are now many researchers using agent-based modelling to concentrate on issues, such as emergence, using agents employing simple heuristics or evolutionary learning algorithms — this area is often termed ‘Agent-based Computational Economics’ (ACE) [16].

We contrast the ‘sociologically inspired’ approach we overview in this paper with a classical game theoretic approach — specifically we are more interested in dynamics than equilibrium and in the development of algorithms that can function in noisy environments with incomplete information.

4 Example: BitTorrent and World War I

A general issue explored by much computational sociological work is that of maximising the collective performance of a group while allowing individual agents reasonable levels of autonomy. In many situations there arises a contradiction between these two aspects. This kind of thing happens in human societies all the time, for example, when someone decides to not to pay on a short train ride (free-ride) or evade tax by not declaring income.

One way to stop these anti-social behaviours is to impose draconian measures via centralised government control — ensuring all individuals behave for the common good stopping free-riders. However, this is costly and hard to police and raises other issues such as: who polices the police? In the parlance of distributed systems engineering — the method does not scale well, is sensitive to noise and has a high computational overhead.

In the context of actually deployed massively distributed software systems, Peer-2-Peer (P2P) file sharing applications (such as the KaZaA and eDonkey systems) have similar problems — most users only download files rather than sharing them [1]. This limits the effectiveness of such systems. Even when the P2P client software is coded to force some level of sharing, users may modify and redistribute a hacked client. It has been noted that P2P file sharing is one of the applications in which only a small number of altruists are needed to support a large number of free riders [1]. Consequently it can be argued that this might be why popular P2P applications tend to be limited to only file sharing rather than, say, processor or distributed storage for example.

These sort of cases can be seen as examples of a more fundamental issue: how can one maintain cooperative (socially beneficial) interactions within an open system under the assumption of high individual (person, agent or peer) autonomy. An archetype of this kind of social dilemma has been developed in the form of a minimal game called the Prisoner's Dilemma (PD) game.

In the PD game two players each selected a move from two alternatives and then the game ends and each player receives a score (or pay-off). Figure 1 shows a so-called 'pay-off matrix' for the game. If both choose the 'cooperate' move then both get a 'reward' — the score R . If both select the 'defect' move they are 'punished' — they get the score P . If one player defects and the other cooperates then the defector gets T (the 'temptation' score), the other getting S (the 'sucker' score). When these pay-offs, which are numbers representing some kind of desirable utility (for example, money), obey the following constraints: $T > R > P > S$ and $2R > T + S$ then we say the game represents a Prisoner's Dilemma (PD). When both players cooperate this represents maximising of the collective good but when one player defects and another cooperates this represents a form of free-riding. The defector gains a higher score (the temptation) at the expense of the co-operator (who then becomes the 'sucker').

A game theoretic analysis drawing on the Nash equilibrium solution concept (as defined by the now famous John Nash [18]) captures the intuition that a utility maximising player would always defect in such games because whatever the other player does a higher score is never attained by choosing to cooperate. The

	Cooperate	Defect
Cooperate	R, R	S, T
Defect	T, S	P, P

Fig. 1. A payoff matrix for the two-player single round Prisoner's Dilemma (PD) game. Given $T > R > P > S \wedge 2R > T + S$ the Nash equilibrium is for both players to select Defect but both selecting Cooperate would produce higher social and individual returns. However, if either player selects Cooperate they are exposed to Defection by their opponent — hence the dilemma

Nash Equilibrium (NE) might be a partial explanation for why there is so much free-riding on existing P2P file-sharing systems — users are simply behaving to maximise their utility. However, do we have any way to solve this problem without going back to centralised control or closed systems? The NE analysis gives us a good explanation for selfish behaviour but not for altruistic behaviour. As stated earlier, even in P2P file sharing systems there are some altruists (keeping the show on the road).

It has been argued by many researchers from the social and life sciences that human societies produce much more cooperation than a Nash analysis would predict. Consequently, various cooperation promoting mechanisms (often using the PD as their test case) have been proposed by social scientists.

BitTorrent, designed by Bram Cohen [4], employs a strategy popularised in the 1980's by computer simulation tournaments applied to the PD. Researchers were asked to submit programs (agents if you like) that repeatedly played the PD against each other [2]. The result of all these tournaments was that a simple strategy called 'Tit-For-Tat' did remarkably well against the majority of other submitted programs.

Tit-for-tat (TFT) operates in environments where the PD is played repeatedly with the same partners for a number of rounds. The basic strategy is simple: an agent starts by cooperating then in subsequent rounds copies the move made in the previous round by its opponent. This means defectors are punished in the future: the strategy relies on future reciprocity. To put it another way, the "shadow" of future interactions motivates cooperative behaviour in the present. In many populations and scenarios this simple strategy can outperform pure defection in the repeated PD.

In the context of BitTorrent, while a file is being downloaded between peers, each peer maintains a rolling average of the download rate from each of the peers it is connected to. It then tries to match it's uploading rate accordingly. If a peer determines that another is not downloading fast enough then it may 'choke' (stop uploading) to that other. Additionally, peers periodically try new peers randomly by uploading to them — testing for better rates [4].

Axelrod used the TFT result to justify sociological hypotheses — such as understanding how fraternisation broke out between enemies across the trenches of World War I. Cohen has applied a modified form of TFT to produce a decen-

tralised file sharing system resistant to free-riding, robust against a number of possible exploitative strategies and scalable.

However, TFT has certain limitations and it is not guaranteed to always be the best way of avoiding free-riding strategies, but its simple to implement and performs ‘well enough’ (currently at least) — BitTorrent traffic currently constitutes a major portion of bandwidth usage on the Internet.

The Tit-For-Tat (TFT) strategy employed by BitTorrent works well when agents exchange many file parts over a period of time (repeat the game interaction many times) but is next to useless if interactions follow a single interaction (such as a single game of the Prisoner’s Dilemma). This tends to limit its use to the sharing of very large files where mutual co-operation can be established.

But how might ”strangers” who interact only once come to co-operate? We discuss a recent technique developed from socially motivated computer models in the next section.

5 Example: File Sharing and the ‘Old School Tie’

Recent work, drawing on agent-based simulations of cooperative group formation based on ‘tags’ (surface features representing social labels or cues [13]) suggests a novel co-operation mechanism which does not require reciprocal arrangements [8, 19]. It is based on the idea of a kind of ‘cultural group selection’ and the well known social psychological phenomena that people tend to favour those believed to be similar to themselves even when this is based on seemingly arbitrary criteria (e.g. wearing the same coloured tie). Like TFT, the mechanism is refreshingly simple. Individuals interact in cliques (subsets of the population sharing the same tags). Periodically, if they find another individual who is getting higher utility than themselves they copy them — changing to their clique and adopting their strategy. Also, periodically, individuals form new cliques and / or randomly change their strategies.

Defectors can do well initially, suckering the co-operators in their clique — but ultimately all the co-operators leave the clique for pastures new — leaving the defectors alone with nobody to free-ride on. Those copying a defector (who does well initially) will also copy their strategy, further reducing the free-riding potential in the clique. So a clique containing any free-riders quickly dissolves but those containing only co-operators grow.

Given an open system of autonomous agents all cliques will eventually be invaded by a free-rider who will exploit and dissolve the clique. However, so long as other new cooperative cliques are being created then co-operation will persist in the population as a whole.

In the sociologically oriented models, cliques are defined as those individuals sharing the same labels and their interpretation is as some kind of socially observable marking attached to individuals. There is no population structure other than the cliques themselves and the population changes over time by employing a population level evolutionary algorithm employing replication and mutation [8, 19].

In the context of application to P2P systems the clique to which a node belongs is defined by its immediate neighbourhood. Movement between cliques and copying of strategies follows a process of network ‘re-wiring’ which brings a form of evolutionary process into the network — an Evolutionary Rewiring Algorithm (ERA). Figure 2 gives an example of this simple re-wiring process followed by each node over time.




Fig. 2. An illustration of ‘replication’ and ‘mutation’ as applied in the Evolutionary Rewiring Algorithm (ERA) from [12]. Shading of nodes represents strategy. In (a) the arrowed link represents a comparison of utility between A and F. Assuming F has higher utility then (b) shows the state of the network after A copies F’s links and strategy and links to F. A possible result of applying mutation to A’s links is shown in (c) and the strategy is mutated in (d).

The adapted tag mechanisms have been shown to be effective in a simulated P2P file-sharing scenario [12] based on that given by Sun et al [22]. The mechanism demonstrates high scalability with zero scaling cost i.e. it does not take longer to establish cooperation in bigger populations (see figure 3). Although there are outstanding issues to be addressed before the technique can be deployed it offers applications beyond file sharing (such as load sharing or co-operative routing). The ERA algorithm bears some comparison with the SLIC algorithm [22] which makes use of incentives. The ERA appears to achieve similar results by producing an emergent incentive structure.

The tag-based process has been likened to ‘old school tie’ in-group effects [20, 9] that appear to permeate many human societies. It offers a possible explanation for why individuals may behave more altruistically towards perceived in-group members, even if they have never met before — a puzzle for self-interest based social theory. Here we have given an overview of how the same mechanism was adapted and applied within a simulated file-sharing P2P scenario to control free-riding when nodes act selfishly [12].




Fig. 3. The chart shows the number of cycles required before high file-sharing behaviour is attained. Ten independent runs for each network size are shown. Note that increasing the network size does not increase the time to high performance — from [12]

6 Prospect: Specialisation with ‘Foraging Tribes’

Specialisation between individuals is the basis of human society. Agents come to specialise in particular tasks and then use methods of exchange or communal ownership to meet the needs of the collective. But how can agents with only local knowledge and simple learning rules come to specialise in this way — particularly if they behave selfishly?

Some models have demonstrated how group processes similar to those discussed previously (i.e. tag-based) can produce internally specialised co-operative groups [10, 11, 21]. Instead of agents evolving behaviours relating to just co-operation or non-co-operation they evolve discreet skill-types in addition to altruistic giving behaviour.

In [10, 11] a resource foraging and harvesting scenario is modelled. Agents forage for resources and then harvest them to gain energy. Different resources require different skills but agents can only possess one skill at a time and are therefore only able to harvest those resources that match their specific skill. An agent may pass a resource it can not harvest to a fellow agent at a cost to itself (an altruistic act) or it may simply ignore such resources (act selfishly). When an agent harvests a resource it attains energy (utility) which can be considered as a form of ‘fitness’. Figure 4 gives a schematic of the scenario.

If agents follow a tag-based evolutionary algorithm (similar to that previously described) then they form groups (which can be thought of as cliques or ‘tribes’) that contain a diversity of skills within them and sharing becomes high.

Figure 5 gives some results from [10]. The main result worth noting is that donation rates are high even when the cost of giving is high to the donating agent. The cost values given are as a proportion of the the harvest value of a resource (one unit of energy).

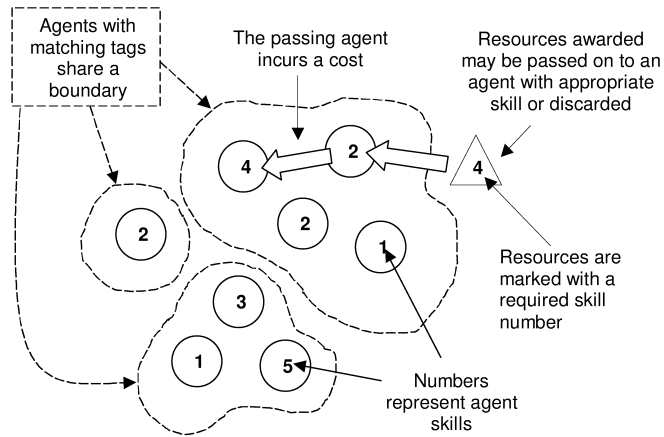


Fig. 4. A schematic representation of how resources are passed to an in-group with the required skill at a cost to the passing agent and hence making use of in-group altruism (from [11]).




Fig. 5. The chart shows averaged results from a number of runs where there are five skills associated with five unique resource types. The x-axis indicates how 'resource rich' the environment is. The y-axis indicates the amount of altruistic donation within groups. The comparison of dumb and smart agents refers to the method of locating a recipient for the donation and the cost indicates the cost to the donating agent (from [10]).

As can be seen, even when donation costs half as much as a harvested resource, donation rates are still high if the environment is sufficiently ‘resource rich’ and a ‘smart’ method of locating recipients is used (the smart method simply means that agents are able to locate others within their group directly rather than search randomly in the population for them — we do not concern ourselves here with this issue).

We can envisage prospects for application of this technique to the formation of internally specialised cliques within P2P networks. The skills would become different kinds of services that nodes could offer (e.g. processing, query answering, storage) and resources could represent job requests submitted at nodes. Figure 6 shows a schematic of this.

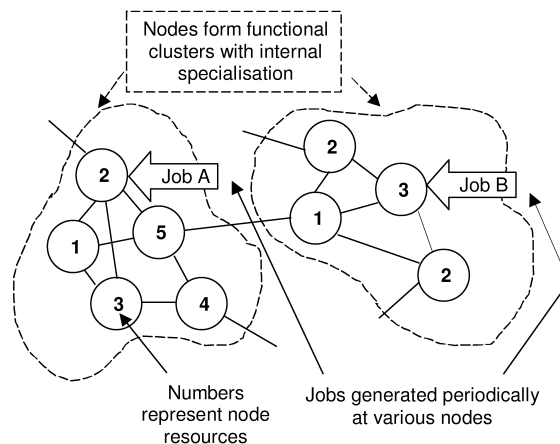


Fig. 6. The specialisation mechanism could be applied within a peer-to-peer network. The above schematic shows an example network fragment. Jobs are submitted at nodes and may require services (or resources) from other nodes. Using a similar mechanism to the ERA algorithm described previously, the network could be made to self-organise into functional clusters to satisfy job requests.

The process of translation from the abstract sociologically oriented models previously produced [10,11] to a P2P type application is a non-trivial exercise — for example, the previous exercise of applying ‘tag’ models of co-operation to P2P file-sharing involved a four stage process in which an abstract model was adapted towards an application domain [12]. At each stage a simulation model needed to be extensively explored to ensure that the desirable emergent properties had not been lost.

However, we are given confidence that specialisation can be generated within working systems since recent work, applied to simulated robotics, applying sim-

ilar techniques based on tags (combined with genetic programming) produced specialised and altruistic behaviour within in-groups (or ‘tribes’) [21].

7 Prospect: Power, Leadership and Hierarchy

A major area of interest to social scientists is the concept of power — what kinds of process can lead to some individuals and groups becoming more powerful than others? Most explanations are tightly related to theories of inequality and economic relationships, hence this is a vast and complex area.

Here we give just a brief very speculative sketch of recent computational work, motivated by sociological questions, that could have significant import into understanding and engineering certain kinds of properties (e.g. in peer-to-peer systems), in which differential power relationships emerge and may, perhaps, be utilised in a functional way.

Interactions in human society are increasing seen as being situated within formal and informal networks [16]. These interactions are often modelled using the abstraction of a game capturing interaction possibilities between linked agents [24]. When agents have the ability to change their networks based on past experience and some goals or predisposition, then, over time, networks evolve and change.

Interestingly, even if agents start with more-or-less equal endowments and freedom to act, and follow the same rules, vastly unequal outcomes can be produced. This can lead to a situation in which some nodes become objectively more powerful than other nodes through topological location (within the evolved network) and exploitative game interactions over time.

Zimmerman et al found this in their simulations of agents playing a version of the Prisoner’s Dilemma on an evolving network [24]. Their motivation and interpretation is socio-economic: agents accumulate ‘wealth’ from the payoffs of playing games with neighbours and make or break connections to neighbours based on a simple satisfaction heuristic (based on a rule discussed in [15]).

Figure 7 (from [24]) shows an example of an emergent stable hierarchical network structure. Interestingly, it was found that, over time, some nodes accumulate large amounts of ‘wealth’ (through exploitative game behaviour) and other nodes become ‘leaders’ by being at the top of a hierarchy. These unequal topological and wealth distributions emerge from simple self-interested behaviour within the network. Essentially, leaders, through their own actions, can re-arrange significantly the topology of the network — those on the bottom of the hierarchy have little ‘topological power’.

The idea of explicitly recognising the possibility of differential power between sub-units in self-* systems and harnessing this is an idea rarely discussed in engineering contexts but could offer new ways to solve difficult co-ordination problems.

Considering P2P applications, one can envisage certain kinds of task in which differential power would be required for efficient operation — e.g. consider two nodes negotiating an exchange on behalf of their ‘group’ or ‘follower’ nodes.




Fig. 7. Forms of ‘hierarchy’, ‘leadership’ and unequal wealth distribution have been observed to emerge in simulated interaction networks (from [24]). Nodes play PD-like games with neighbours and break connections based on a simple satisfaction rule. Hierarchies are produced in which some nodes are more connected and hence can effect the network dramatically by their individual actions — a form of ‘topological power’.

This might be more efficient than individual nodes having to negotiate with each other every time they wished to interact. Or consider a node reducing intra-group conflict by imposing a central plan of action.

We mention the notion of engineering emergent power structures, briefly and speculatively here, because we consider power to be an under-explored phenomena within evolving information systems. Agents, units or nodes are often assumed to have equal power. It is rare for human societies to possess such egalitarian properties and perhaps many self-* like properties are facilitated by the application of unequal power relationships. We consider this a fascinating area for future work.

8 Conclusion and Summary

Here we have provided some examples and prospects of sociologically inspired approaches to engineering self-* systems. Rather than attempt an extensive overview we have focused on a few encouraging specific results and possible P2P-type applications.

We believe that the computational social science literature can be a potential source of new techniques and ideas for prospective self-* engineer because social phenomena are generally self-organising, robust and scalable — all desirable properties for self-organising information systems.

Computational social science tries to reverse engineer general properties at a fairly abstract level whereas self-* engineers need to apply techniques to specific concrete problem domains. As we have hoped to show, however, it is possible to import useful techniques (see [12] for a case study in applying a technique to realistic domain) from the one approach to the other.

The idea of using social metaphors and approaches for the construction of smart information systems is far from new [17]. What is new is that distributed systems engineers are increasingly asking sociological questions (even if they are unaware of it!) and social scientists are increasingly turning to algorithmic specification and computer simulation to explore their theories. We hope that advances from both areas can be brought together and used to reinforce each other. Experience so far indicates this not to be an unreasonable hope.

Acknowledgements

This work partially supported by the EU within the 6th Framework Program under contract 001907 (DELIS).

References

1. Adar, E. and Huberman, B.: Free Riding on Gnutella. *First Monday* Volume 5, No. 10. (2000).
2. Axelrod, R.: *The evolution of cooperation*. N.Y.: Basic Books, (1984).
3. Binmore, K.: *Game Theory and the Social Contract. Volume 2: Just Playing*. Cambridge, MA: The MIT Press, (1998).
4. Cohen, B.: Incentives Build Robustness in BitTorrent. Presented at the *1st Workshop on the Economics of Peer-2-Peer Systems*, June 5-6, 2003, Berkley, CA. Available at: <http://www.sims.berkeley.edu/research/conferences/p2pecon/>
5. Epstein, J.M. and Axtell, R.: *Growing Artificial Societies: Social Science From The Bottom Up*. London: MIT Press, (1996).
6. Gilbert, N. and Doran J., (eds.): *Simulating Societies: the Computer Simulation of Social Phenomena*. London: UCL Press, (1994).
7. Gilbert, N. and Conte, R. (eds.): *Artificial Societies: the Computer Simulation of Social Life*. London: UCL Press, (1995).
8. Hales, D.: Cooperation without Space or Memory: Tags, Groups and the Prisoner's Dilemma. In Moss, S., Davidsson, P. (eds.) *Multi-Agent-Based Simulation. Lecture Notes in Artificial Intelligence 1979*. Berlin: Springer-Verlag, (2000).
9. Hales, D.: Tag Based Cooperation in Artificial Societies. Ph.D. Thesis, Department of Computer Science, University of Essex, UK, (2001).
10. Hales, D.: Evolving Specialisation, Altruism and Group-Level Optimisation Using Tags. In Sichman, J. S., Bousquet, F. Davidsson, P. (eds.) *Multi-Agent-Based Simulation II. Lecture Notes in Artificial Intelligence 2581*. Berlin: Springer-Verlag, (2002).
11. Hales, D.: Searching for a Soulmate — Searching for Tag-Similar Partners Evolves and Supports Specialization in Groups. In Lindemann, G., Moldt, D. and Paolucci, M., (eds.) *Regulated Agent-Based Social Systems — 1st International Workshop, Lecture Notes in Artificial Intelligence 2934*. Berlin: Springer-Verlag, (2004).

12. Hales, D. (2004) From selfish nodes to cooperative networks — emergent link based incentives in peer-to-peer networks. In *Proc. of the 4th IEEE International Conference on Peer-to-Peer Computing (P2P2004)*. IEEE Computer Soc. Press, (2004).
13. Holland, J.: The Effect of Labels (Tags) on Social Interactions. Santa Fe Institute Working Paper 93-10-064. Santa Fe, NM, (1993).
14. The Journal of Artificial Societies and Social Simulation (JASSS). Available at: <http://jasss.soc.surrey.ac.uk>
15. Kirman, A.: Ants, Rationality and Recruitment. *Quarterly Journal of Economics*, 108, 137156, (1993).
16. Kirman, A.P., and Vriend, N.J.: Evolving Market Structure: An ACE Model of Price Dispersion and Loyalty. *Journal of Economic Dynamics and Control*, 25, Nos. 3/4, 459-502, (2001).
17. Minsky, M.: *Society of Mind*. Simon & Schuster, (1988).
18. Nash, J. F.: Equilibrium Points in N-Person Games, *Proc. Natl. Acad. Sci. USA* 36, 48-49, (1950).
19. Riolo, R., Cohen, M. D. & Axelrod, R.: Cooperation without Reciprocity. *Nature* 414, 441-443, (2001).
20. Sigmund & Nowak: Tides of tolerance. *Nature* 414, 403-405, (2001).
21. Spector, L., J. Klein, C. Perry, and M. Feinstein.: Emergence of Collective Behavior in Evolving Populations of Flying Agents. In E. Cantu-Paz, et al (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2003)*, pp. 6173. Berlin: Springer-Verlag, (2003).
22. Sun Q. & Garcia-Molina, H.: SLIC: A Selfish Link-based Incentive Mechanism for Unstructured Peer-to-Peer Networks. In *Proceedings of the 24th IEEE international Conference on Distributed Systems*. IEEE computer Society, (2004).
23. von Neumann, J. and Morgenstern, O.: *Theory of Games and Economic Behavior*. Princeton, (1944).
24. Zimmermann, M.G., Egufluz, V.M. and San Miguel.: Cooperation, adaptation and the emergence of leadership. In A. Kirman and J.B. Zimmermann (eds.) *Economics with Heterogeneous Interacting Agents*, pp. 73-86. Berlin: Springer, (2001).