Tag Mechanisms Applied to Open Grid Virtual Organizations Management

Isaac Chao, Oscar Ardaiz and Ramon Sanguesa
Software Department, Polytechnic University of Catalonia, Spain
{ichao,sanguesa}@lsi.upc.es, oscar.ardaiz@unavarra.es

Abstract

A key challenge in Grid computing research is the design of efficient and self-organizing management mechanisms of the system's Virtual Organizations (VOs). The larger the scale, heterogeneity and dynamism of the target Grid, the more important becomes the inclusion of self- properties. In this paper we explore the application of novel self-organizing techniques to the management of open VOs systems. We face the problem of emergence and cooperation establishment between self-interested agents interacting in a highly dynamic and heterogeneous ecosystem of VOs. We will show how a socially-inspired technique based on Tags can be used as the mechanism leading to the emergence of cooperation in such systems. The contribution of this paper is matching the cooperative groups emerging from Tag mechanisms with the VOs in the Grid; this allows addressing important open issues from previous Tag models by introducing complementary management mechanisms at the VO level. This extension to previous Tag models allows the system to cope with the presence of different types of non-adaptive agents. We provide a complete model and we evaluate its performance with experiments in a simulator.*

1. Introduction

Virtual Organizations (VO) are geographically distributed, functionally diverse, dynamic and agile organizational entities linked through Information and Communications Technology (ICT). In recent years Grid computing evolved from sharing computing power and scientific instruments between supercomputers to the sharing and exchange between customers, suppliers and partners of generic Grid services structured in VOs [JEF04]. Next generation Grids focus will be knowledge sharing to enabling collaborations between different VOs while respecting their individual policies [EGR06]. We target the Inter-Grid scenario depicted in Figure 1. These systems of VOs are often composed of very heterogeneous components, showing diverse interests and belonging to different administrative domains. The following characteristics are probable:

- Lack of previous interactions between agents or any usable historical data
- High dynamicity: Agents entering and exiting the system continuously
- Impossibility of using a centralized institution to enforce fulfilment of agreements
- Ad-hoc grids can be formed by communities that continuously change their usage policies, membership and goals during the lifetime of the Grid [ALM04]

These kinds of systems fall into the so called Decentralized Autonomic Computing [WoHo06]. A centralized control or manual management is exceedingly difficult or even impossible. Decentralized, self-organized management and operation becomes a requirement. Two important challenges arise: Firstly the system must be engineered to cope with high levels of heterogeneity and dynamicity. We propose a self-organizing mechanism based on Tags to coordinate VOs evolution. This shows emergence of the required macroscopic properties from local agent's interactions. Second, the engineered protocols must be robust in the presence of self-interested agents, capable to operate in pursuit of local interests, even at the expense of global system objectives. The model we propose provides the agents with the right incentives to promote cooperation, achieving global system objectives rather than individual self-interested local objectives, even in the presence of non-adaptive agents.

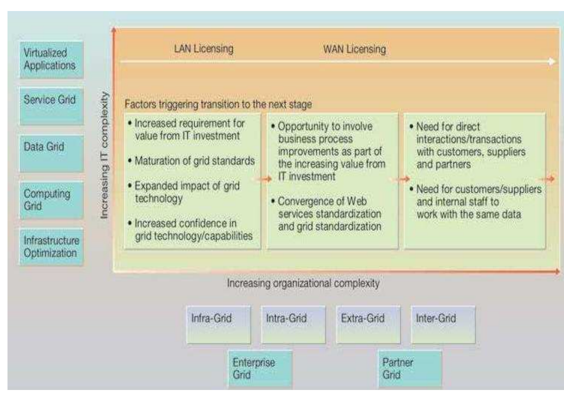

Figure 1: Evolution of Grid Computing (from [JEF04])

The rest of the paper is structured as follows. In section 2 we present the context of the target domain and related work. In section 3 we discuss state-of-the-art Tag mechanisms and we evaluate their limitations. In section 4 we describe our Tag mechanism for VOs. In section 5 we present experimental results and evaluation. Section 6 concludes the paper.

2. Related Work

Virtual Organizations (VOs) for the Grid as introduced in [FKT01] require very flexible sharing relationships. They are formed with the goal of performing resource sharing and coordinated problem-solving in dynamic, multi-institutional environments. VOs can be created to undertake their role for a very brief period of time, or exist for a longer term. They may be created on demand in dynamic, open and competitive environments. This includes the scenarios targeted in decentralized Grid markets [Catn04, ACC+05]. The basic operations for a VO are its creation, formation, operation and evolution [CaAf03]. The automated management of these operations is still a research challenge.

There are many past and ongoing projects using VOs conceptually, but very few projects are addressing the management of VOs themselves. In Conoise-G project [PTL+05] an agent system supporting robust and resilient VOs formation and operation is presented. Another project focusing on Trust issues is Tustcom [TCom05] aiming to provide a trust and contract management framework enabling the definition and secure enactment of collaborative business processes within VOs that are formed on-demand, self-managing and evolve dynamically. In both Conoise-G and Trustcom approaches to VO management, components for helping automated VO management are developed, but no specific self-organization mechanism is provided. Self-organization mechanisms incorporating emergence bring into VO management higher levels of flexibility and adaptability, providing much more generic models, applicable for a wider range of scenarios.

In this paper we focus on the issue of tracking non-adaptive agents exhibiting different cheating strategies. This sub-domain of VO management can be linked to reputation management. A complete and updated survey on the most salient challenges in this domain can be found in [Dell05]. In the specific grid VOs domain, the most relevant work has been carried out in the GridEigenTrust framework [KVLA03]. This work presents a complete architecture for a reputation management service across different VOs and contexts. This is an example a social networks based reputation mechanism. However, the model has an important drawback since it relies on a set of pre-trusted peers in order to inject trust on the system. For agents entering in an un-trusted open system composed of VOs it might be impossible to agree on a set of common pre-trusted peers.

3. Tag mechanisms and models

3.1 Tag models background

Holland [Holl93] first proposed the concept of tags as markings or social cues that are attached to agents and are observable by others. Since then, a numbers of tag models have emerged to improve on the classical individual interest versus global social welfare Prisoner's Dilema (PD). Riolo [Riol00] has described a number of tagging approaches to address the iterated PD. These approaches outline basic forms of tagging: fixed-bias tagging, variable-bias tagging and evolved-bias tagging. Hales addresses single round PD [Hale00], where there is no even "shadow of the future" concept to support strategies like tit-for-tat as it happens in the iterated PD [Axel81]. Tags promote the emergence of cooperation between agents. These techniques are attractive since they don't require centralized or third party reputation systems, the monitoring of neighbour behaviour or the explicit programming of incentives. They also can be used in highly dynamic environments. These techniques have already successfully been applied to P2P scenarios [Hales04], though showing several limitations (see section 3.2)

In Hales Tag model, each agent is represented by a small string of bits. Interaction involves pairs of agents playing a single round of PD. Agent bits are initialized at random. One bit is designated as the PD strategy bit: agents possessing a "1" bit play cooperate but those possessing a "0" bit play defeat. The other bits represent the agent tag. These bits that have no direct effect on the PD strategy selected by the agent but they are observable by all other agents. In this setting, a very simple algorithm is applied trough a number of rounds: First agents play preferentially with other agents sharing the same tag. Then agents evolve following an evolutionary algorithm which reproduces agent's strategies having collected bigger payoffs. Mutation factors on both tag and strategy are applied. The evolution of the population precipitates a kind of "group selection" process in which those groups (each group being defined by a tag) which are more cooperative tend to predominate but still die out as they are invaded by non cooperative agents. By constantly changing tag strings (by reproduction of those with higher utility) the agents produce a dynamic process that leads to high levels of cooperative actions. Extensive experimentation varying a number of parameters showed that for a big enough Tag space, high levels of cooperation quickly predominated in the population.

3.2 Tag models open issues

Important open issues still need to be addressed to allow for usage of these mechanisms to enforce cooperation in realistic open Grid environments. First and foremost, rational agents with incentives to misreport their utility to others or not copying cooperative strategies from others may decide to never adapt to the mechanism and just move from group to group free-riding. Adapting to the mechanism is a requirement for the agents in state-of-the-art Tag models. This assumption would be clearly unrealistic for our target scenario of VOs

composed of competitive rational agents. The only way to face the moral hazard situation is introducing sanctioning mechanisms [Dell05].
Additionally, as in most evolutionary models, it is assumed that a fitness or utility measure can be always easily extracted for the agent. However this might not be the case for complex domains, as for example a complex problem solving scenario requiring the composition of Grid services workflows over different contexts, with complex dependencies among services. Finally, it is not clear to what extend a learning mechanism based just in copying others behaviours is best suited to a Grid scenario. It might be the case for a Tag model application that just mimicry is not a good option for the system performance [DoSa05]. Even when mimicry is good for the social welfare, in games more complex than PD it might be difficult for an agent to estimate from an observed behaviour the internal strategy run by another agent. An additional component of control at the VO level would be desirable to provide convenient support services to agents and to enforce specific polices for specific tasks and goals in VO.

4. Tag Mechanisms applied to VOs

4.1 The generic model

We aim to extend the described Tag mechanisms with the purpose of addressing its current limitations for a VO scenario, providing robustness to non-adaptive agents as well as complementary supporting services for the agents. In a Grid composed of VOs scenario, agent's interaction is constrained to the VO scope. In order to interact with agents from another VO, an agent needs to enter that VO first. This structuring in VOs is an important difference between Grid systems and pure P2P systems. However this does not necessarily imply a hierarchical structure. The VO functions can be implemented by the decentralized coordination of the set of agents composing the VO. Our target scenario has the following requirements and limitations:

- The Grid is formed by a set of VOs being dynamically composed (VO ecosystem)
- Hierarchical or centralized Grid management solutions are no applicable due to the characteristics of open Grid systems cited in section 1
- The goal is the achievement of high cooperation between agents leading to system-wide overall utility (social welfare) maximization, even in the presence of a sensible number of non-adaptive agents. However collusive attacks are not considered.

A static view of the model is shown in Figure 2. We can identify the agents and the VOs they conform. There are potentially many different types of VOs incorporating various resource management mechanisms, as correspond to a heterogeneous inter-Grid. In order to increase comprehension of the model, we have instantiated these concepts in a Grid Market scenario, where each VO establishes a market in order to allocate resources. Such a scenario can be implemented using state-of-the-art economic Grid middleware such as the proposed in [ACC+05].

The model builds on the fact that groups of agents (sharing a Tag) emerge in Tag models. We make explicit the autonomy of such groups by identifying them with the VOs. We rely on the fact that open Grid systems already incorporate such a concept of VOs as entities with a purpose. Agents operations are strictly constrained to the VO they inhabit, any interaction outside the VO boundary is just for the evolution phase, when learning about the performance of other agents in different VOs. The explicit separation in VOs makes it possible to address aforementioned Tag mechanisms limitations, by means of a two-tier system dynamics: Internal VO dynamics is coordinated by the intra-VO mechanism (incorporating a sanctioning mechanism), whereas VOs co-evolution in the system is driven by the Tag mechanism.

The scope of the generic model is indented to address a wide range of management mechanisms. The model does not imply central coordination at any level. Each VO (understood as the set of agents composing it) maintains autonomy to decide its own coordination mechanism. Each individual agent is still totally free to choose strategy in order to maximize its individual utility and is totally free to move to other VOs. The agents composing each VO at a given moment are able to decide resource allocation mechanisms and to fix any management entities if required (as shown in the diagram). Fully decentralized management can be also agreed between the agents composing the VO for the intra-VO coordination mechanism. Various types of sanctioning mechanism can be applied inside the VO, ranging from centralized blackboards to fully decentralized mechanisms (see for example [PaSt05]).

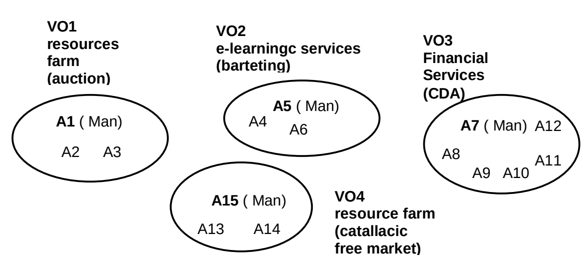

Figure 2: Model's static view

Having the VO as basic unit for control makes it possible to provide support services to model utility comparison in complex interactions (using for example shared ontologies during evolution phase). Additionally, the "blind-copy" evolutionary learning and inter-VO natural selection process might be complemented at the intra-VO level by enforcing specific policies, enabling to meet more realistic VO requirements. The model's algorithm (figure 3) covers the VO lifecycle phases of operation and evolution. Creation and formation phases are considered at the system bootstrapping (see section 5.1). Creation of new VOs eventually happens as a consequence of a Tag mutation.

For a number of rounds
begin main loop
begin interaction phase
Interaction (VO operation phase):
For each VO in the system:
For each agent in the VO:
Interact with another agent from the VO.
Payoffs are generated from interaction
Sanctioning mechanism is applied
end interaction phase
begin evolution phase
Evolution (VO evolution phase):
For each agent in the system:
Reproduce: contact an agent from another VO
Compare agent's payoff from last interaction
If (the other agent outperforms current agent) then
Current agent moves to the other's VO
Agent adapts strategy following its type
else current agent stays on its current VO
Mutate: Agent applies probabilistic Tag mutation
end evolution phase
end main loop

Figure 3: The algorithm

For the interaction phase, each agent starts an interaction with another agent from the same VO. This will normally imply the request, exchange or trading of one or more Grid services. After interaction is accomplished, sanctioning mechanism currently active in the VO is applied when required. For the evolution phase, in the case the agent is adaptive it will also adapt strategy by copying the other agent strategy. In contrary case, it will update its strategy following its specific behavioural rule. The mutation on Tag implies the agent abandoning its current VO and starting a new one on its own. For non-adaptive agents, mutation implies moving to another VO in spite of reproduction phase results. The idea behind Tag mutation is to introduce variability on the system, so it can recover from predominantly defective populations.

A dynamic view of the model is shown in figure 4. After the interaction phase, which may consist from a very simple abstraction of a competitive game to more complicated negotiations implied in a Grid market, it takes place the evolution phase. In figure 4 we can see how agent A7 from VO3 learns that payoffs for agent A6 in VO2 are bigger, and decides to move from VO3 to VO2. When getting into VO2, agent A7 pays a corresponding entrance fee. If A7 behaves cooperatively inside VO2 following the Tag mechanism it will help increasing the size of the cooperative cluster. If A7 is a non-adaptive agent and misbehaves regardless of the outcome of the Tag mechanism trying to free-ride, then the sanctioning mechanism implemented in VO2 must punish A7, providing the agent with incentives to leave the cooperative VO. As for VO3, disbanding phase has arrived and A8 will probably need to look for new partners or join another existing VO.

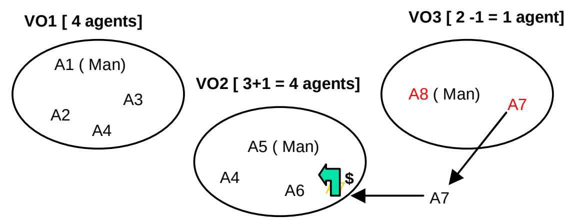

Figure 4: Model's dynamic view: Snapshot at Evolution Phase

4.2 Implementation of VOWorld

In a first implementation of the generic model (VO World) we designate a VO Manager for each VO, and a blackboard controlled by the VO Manager on each VO is used as sanctioning mechanism. The blackboard keeps records of agent's misbehaviours within the VO, and helps adaptive agents to filter their partner agents for interaction within the VO in order to minimize free-riding. The blackboard consists of two variables: a sanction period for the agent within the VO (the sanctioned agent will not be able to interact within that VO for that period), and a maximum allowed number of total defeat operations for any agent in the VO, for the VO lifetime (this models the maximum level of free-riding allowed for a single agent on a VO). During the sanction period and resulting from lower payoffs obtained, sanctioned agents might migrate to other VOs were they enrol in interactions again. We found 5 rounds and 5 maximum defeats as good performing values for the blackboard. However a wide range of other values would make also a working blackboard for this scenario, but slightly less performing. Specifically, the maximum allowed number of total defeats per agent on the VO can be set up to a much exaggerated value of 1000 and still the performance would be comparable with the best performing case, and clearly outperforming the case without blackboard.

Regardless of which shape the intra-VO coordination mechanism takes, it is important to notice that the inter-VO Tag mechanism greatly impacts the intra-VO mechanism. As we have tested in our current implementation of the model, a centralized intra-VO management by a VO manager gets its behaviour affected by the interleaving with the inter-VO Tag mechanism (which is coordinating at each time step which agents belong to which VO). There is certainly transference of self-organization from the inter-VO mechanism to the intra-VO mechanism. A more advanced interesting effect to study would be the interaction between two self-organizing mechanism, by including an appropriate mechanism (maybe based also on Tags) for the intra-VO management.

We also make two plausible economical assumptions for this concrete implementation of the model: First, adopting the VO manager role has a cost; this assumption tries to prevent non-adaptive agents to become VO managers. This role is typically reserved to altruistic agents. This cost variable can be tuned (but we are not doing so here) to provide non-adaptive agents with the right incentives not to become VO Managers. Second, some cost is also implied on moving from one VO to another. This addresses cheap pseudonyms issue, critical in decentralized reputation mechanism design.

For the modelling of the interaction between agents we use in this implementation the single round PD game. The PD is very simple but captures the basic characteristics from a competitive exchange game incorporating social dilemma. We use the same PD payoff matrix from Hales Tag models: A reward payoff (R) and a punishment payoff (P) are given for mutual cooperation and mutual defection respectively. When different moves are selected, differential payoffs of temptation (T) and sucker (S) are awarded to defector and co-operator respectively. $T > R > P > S$, and the constraint $2R > T + S$ makes a PD. We set T=1900, R =1000 and P=S=1.

5. Experimental results evaluation

5.1 Simulation environment

The simulation performs an initial bootstrapping phase, distributing agents in several VOs at start-up, and a VO manager is designated for each VO. Experimental results shown in section 5.2 are consistent in the ranges:

0 .001 < tm < 0.1, tm is the tag mutation
10 < N < 10000, N is the number of agents
N/10< nV <N/2, nV is the initial number of VOs

The results are thus valid for large scale scenarios, allow for a wide screen on mutation levels to be selected, and are almost insensitive to the bootstrapping of the agents in VOs at start-up. Flexibility in mechanism parameters choice translates into flexibility when engineering real open Grid systems with the described model.

5.2 Experiments with non adaptive agents

We evaluate the proposed mechanism robustness to different non-adaptive agents. We fix Tag mutation in 0.01 and distribute a population of 100 agents in an initial bootstrapping of 50 VOs, designating one manager per VO. A fraction (10%) of different types of non-adaptive agents is injected in the initial population. We measure the cooperation level in each round as the proportion of agents in the population playing cooperate in the PD game (Figure 5).

Types of cheaters:

-*Pure defective*: Always play defeat, trying to free-ride from cooperative agents

-*Nihilistic*: Always play defeat, and always report have played defeat, so other agents become cheaters too. The target of nihilistic agents is to destroy the mechanism itself by spreading defection

-*Mixed-Probabilistic*: They play cooperate or defeat in a probabilistic basis. These can be considered a kind of outsiders who don't follow clear incentives.

-*Tit-for-tat*: Apply tit-for-tat strategy; they play the same action as the last partner they encountered. Not to mistake with the tit-for-that in iterated PD. Here each partner is a different agent.

-*Greedy*: Always play defeat, and always report having played cooperative (so other agents will keep cooperating and they can continue free-riding on them)

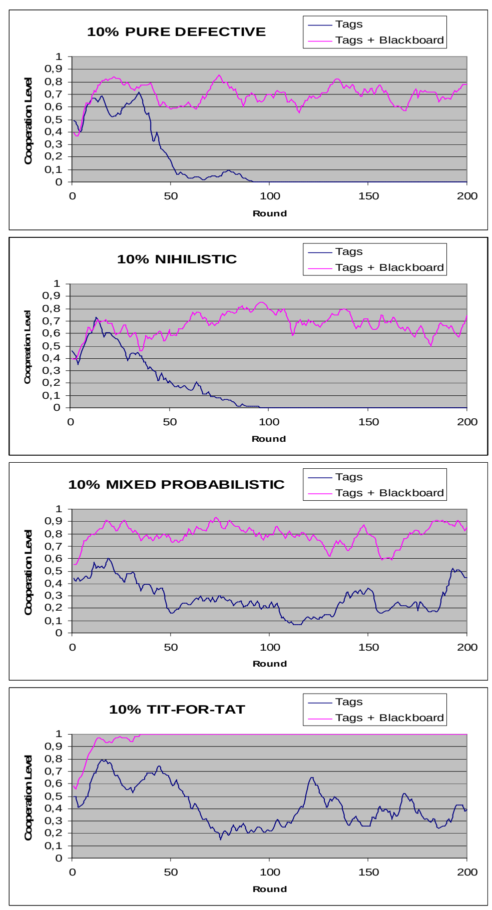

Figure 5: Experimental results with non-adaptive agents

The results in graphs from figure 5 show how, with pure defective agents using just Tags, the population quickly gets infected by free-riding and it is not able to recover in the long term. Contrarily, the system with blackboards constantly adapts its VOs to the presence of defectors, by sanctioning them, achieving on its peaks a cooperation level close to the theoretical maximum (for a 10% of pure defectors this is 0.9). For the nihilist population we can see that without sanctioning those agents are able to totally destroy cooperation even faster than for pure defectors, but the blackboard is able to cope with them and stabilize on an acceptable level of cooperation. Mixed probabilistic populations are easily controlled by the blackboard, but not without it. Tit-for-tat strategies are greatly managed by the blackboard, but the case without blackboard slowly derives to a total defective population.

For greedy agents (Figure 6) total cooperation level is quickly achieved in both cases, very close to what corresponds to a pure adaptive population (hence graph is not shown). It would appear that there is no gain by using the blackboard. However greedy agents are not targeting to destroy cooperation, but to promote it in order to better free-ride from others. It is important to evaluate to what extent they are achieving their goal. Interestingly, the average utility of Greedy agents in each round is much lower when using blackboard. Without blackboard they are almost constantly getting the temptation payoff of the PD game (1900). With blackboard, their utilities get much decreased, even lower than the mutual reward payoff (1000), and hence for them cheating become even worst in performance than cooperating in a fully cooperative VO. This means the blackboard dramatically reduces the level of free-riding non-adaptive agents are able to commit at expenses of adaptive agents utilities.

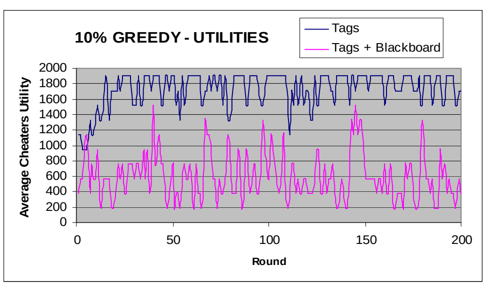

Figure 6: Utility for non-adaptive agents

For all kind of cheaters evaluated, we have considered also non adaptive agent's misreporting their utilities during evolution phase. Cheaters might have incentives to advertise themselves as good performing in order to better free-ride from agents moving to their VOs. Experimental results obtained were similar to the ones showed here. The explanation for this is that the system dynamics constantly couples both inter-VO mechanism (Tags) whit intra-VO sanctioning (in this case a blackboard). Agent's been exploited are always aware through the Tag mechanism of how much better they could be performing in other VOs (and consequently move to those VOs) and cooperative VOs can protect themselves from free-riding agents invading them relying on the intra-VO sanctioning mechanism. A metaphor commonly used to exaplain Tag mechanism is social groups or tribes differentiated by some observable trait (e.g. clothes in urban tribes or skin colour in ghettos) which is used to build ad-hoc trust between individuals sharing the trait [Hale00]. A corresponding metaphor for the additional intra-VO sanctioning mechanism would be taking into account the enforcement of the specific behavioural code of the tribe. Showing no adaptation to this code might result in individual loosing its initial reputation and having to move to other tribes or upon rejection in all other tribes become isolated or relegated to share with other non-adapted individuals.

6. Conclusion

We have presented a novel approach to engineer VOs in open Grid systems based on Tag mechanism, achieving cooperative behaviours as an emergent property resulting from the group dynamics and VO "natural selection" at the inter-VO. The system self-organizes achieving the emergence of cooperation-based VOs without the intervention of any component enforcing that property explicitly across VOs. The model incorporates complementary sanctioning mechanisms at the intra-VO level, enabling the system to cope with various types of non-adaptive agents. This intra-VO mechanism builds on the very nature of a system composed of VOs (prototypically a Grid). Hence we do not impose any coordination or management by exogenous components. The intra-VO mechanism enables to overcome important open issues in state-of-the-art Tag models. Experimental results show high improvement in system wide cooperation level for the cases where the intra-VO sanctioning mechanism is used in the presence of non-adaptive agents, as well as an important decrease of free-riding these agents are able to commit at expenses of the overall system utility. This makes the model more robust and applicable to more realistic VO scenarios.

Future work will consist on the exploration of variations on the Tag mechanisms (allowing for agents to belong to various VOs simultaneously) and the use of alternative sanctioning mechanisms within the VO. A parametric study of relevant macroscopic variables (VO sizes, VO stability, etc) will be accomplished from microscopic mechanism variables tuning. Modelling an application scenario different from the PD and closer to a grid VO will provide a bigger set of microscopic parameters to experiment with towards the emergent and self-organized engineering of grid VOs using Tag mechanisms.

Acknowledgements

This work was supported in part by European Union under Contract CATNETS EU IST-FP6-003769 and the Spanish Government under Contract TIC2002-04258-C03-01

7. References

[ACC+05] Oscar Ardaiz, Pablo Chacin, Isaac Chao, Felix Freitag, Leandro Navarro, An Architecture for Incorporating Decentralized Economic Models in Application Layer Networks, in Proceedings of the 1st Int. Workshop on Smart Grid Technologies, Utrecht, Netherlands, July 2005

[ALM04] Toward an Architecture for Ad Hoc Grids. Kaizar Amin, Gregor von Laszewski, and Armin R. Mikler. Proceedings of the IEEE 12th International Conference on Advanced Computing and Communications (ADCOM 2004). December 2004, Ahmedabad, India.

[Axel81] The Evolution of Cooperation Axelrod, Robert. Science, 211(4489):1390-6 (1981).

[CaAf03] Camarinha-Matos, LM and Afsarmanesh, H. (2003) A Roadmap for Strategic Research on. Virtual Organisations, in Proceedings of PRO-VE 2003, 33-46, Kluwer

[Catn04] CATNETS Project (2004), “Annex I – Description of work”, IST-FP6-003769

[Dell05] Chrysanthos Dellarocas, Reputation Mechanisms, Handbook on Economics and Information Systems, June 2005

[DoSa05] Austin McDonald and Sandip Sen, "The Success and Failure of Tag-Mediated Evolution of Cooperation," in the Working Notes of the AAMAS-05 Workshop on Learning and Adaptation in MAS (LAMAS), 2005

[EGR06] “Next Generation Grid(s) - European Grid Research 2005 - 2010”, June 2003

[FKT01] I. Foster, C. Kesselman, S. Tuecke. The Anatomy of the Grid: Enabling Scalable Virtual Organizations. International J. Supercomputer Applications, 15(3), 2001.

[Hale00] Hales, D. (2000) 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 Springer-Verlag

[Hales04] Hales, D., From Selfish Nodes to Cooperative Networks - Emergent Link-based Incentives in Peer-to-Peer Networks. Proceedings of The Fourth IEEE International Conference on Peer-to-Peer Computing (p2p2004), Switzerland, 2004, IEEE Computer Society Press

[Holl93] J. Holland. The effects of labels (tags) on social interactions. Working Paper Santa Fe Institute 93-10-064, 1993.

[JEF04] "Evolution of grid computing architecture and grid adoption models," J. Joseph, M. Ernest, and C. Fellenstein, IBM Press, 2004

[KVLA03] Beulah Alunkal,Ivana Valjkovic, Gregor von Laszewski, and Kaizar Amin Reputation-based Grid Resource selection. Workshop of Adap- tive Grid Middleware, New Orleans, LA, September 2003

[PaSt05] Th. G. Papaioannou and G. D. Stamoulis. An Incentives' Mechanism Promoting Truthful Feedback in Peer-to-Peer Systems. In Proc. of IEEE/ACM CCGRID 2005 (Workshop on Global P2P Computing), May 2005

[PTL+05] J. Patel, L. Teacy, M. Luck, N. R. Jennings, S. Chalmers, N. Oren, T. J. Norman, A. Preece, P. M. D. Gray, P. J. Stockreisser, G. Shercliff, J. Shao, W. A. Gray, N. J. Fiddian, and S. Thompson Agent-based virtual organisations for the grid, in Proceedings of the 1st Int. Workshop on Smart Grid Technologies, Utrecht, Netherlands, July 2005

[Riol00] R. Riolo. The efects of tag-mediated selection of partners in evolving populations playing the iterated prisoners dilemma. Nature 414, pages 441–443, 2000.

[TCom05] http://www.eu-trustcom.com/

[WoHo06] T. De Wolf and T. Holvoet. A Taxonomy for Self-* Properties in Decentralised Autonomic Computing, Chapter in 'Autonomic Computing: Concepts, Infrastructure, and Applications', M. Parashar and S. Hariri (editors), to appear in august 2006