

Tag-Based Cooperation in Peer-to-Peer Networks with Newscast¹

Andrea Marcozzi ^a, David Hales ^a, Gian Paolo Jesi ^a, Stefano Arteconi ^a,
Ozalp Babaoglu ^a

^a *University of Bologna Dept. of Computer Science*
Via Mura Anteo Zamboni, 7
40127 Bologna, Italy
{marcozzi, hales, jesi, arteconi, babaoglu}@cs.unibo.it

Abstract. Recent work has proposed how socially inspired mechanisms, based on “tags” and developed within social science simulations, might be applied in peer-to-peer overlay networks to maintain high cooperation between peers even when they act selfishly. The proposed mechanism involves a dynamic re-wiring algorithm called “SLAC”. The algorithm assumes a random sampling service over the entire population of nodes but does not implement this itself. In this paper we re-implement SLAC on an open source peer-to-peer simulation testbed called “PEERSIM”. For the random sampling service we utilize an existing protocol called “NEWSCAST”. We present the results of some experiments we performed in which peers play the Prisoner’s Dilemma game with their neighbours. Our results demonstrate that SLAC augmented with NEWSCAST produces high levels of cooperation. This increases our confidence that previous results from SLAC are generally applicable and valid and also that SLAC could have applications in real implemented systems. Finally we discuss the open issues that need to be addressed for SLAC to progress to a valuable deployable protocol.

Keywords. peer-to-peer, cooperation, social networks, adaptivity, tags

1. Introduction

In recent works a novel socially inspired algorithm based on the “tag” idea [7] has been applied to the problem of sustaining cooperation in peer-to-peer networks composed of nodes behaving selfishly [4].

In the earlier tag models, individual agents interact randomly (in the form of mean-field interaction) under the constraint that they are more likely to interact with agents sharing an identical or similar tag.

In the context of the social scientific interpretation tags represent arbitrary surface markings attached to agents that can be observed and copied by other agents. They have no direct behavioural significance [7]. In human society tags can be viewed as fashions:

¹Based on the Laurea thesis of Andrea Marcozzi, March 2005. Partially supported by the EU within the 6th Framework Programme under contract 001907 “Dynamically Evolving, Large Scale Information Systems” (DELIS).

styles of dress, colors of hat, brand logos or cosmetic makeup. The key property of tags is that although they are distinctive and immediately observable they can be quickly changed and copied.

A number of simulation models have demonstrated that over a broad range of parameter values high levels of cooperation and altruism emerge when agents bias their interaction towards others sharing the same tag [5,11]. These previous models follow an evolutionary approach in which agents reproduce and mutate their behaviours and tags. In such models those gaining higher utility are more likely to reproduce their tags and behaviours with the addition of small amounts of mutation - noise or random variation. The assumptions underlying this approach is that such an evolutionary mechanism can capture the essential elements of cultural learning within a society - that agents are bounded optimizers, copying the tags and behaviours of those who are gaining higher utility than themselves.

We have previously described, in detail, the modification needed to transform (via a series of simulation models) the evolutionary models into an algorithm more applicable to a target application of peer-to-peer file sharing [4]. However, in that work the final algorithm (the SLAC algorithm) was tested on a non-publicly available and highly abstracted simulation test bed. Also, SLAC relies on a random sampling service over the entire P2P network and this was assumed rather than implemented.

In this paper we describe a re-implementation of, and experiments with, SLAC on the open source PEERSIM system [13]. PEERSIM offers a more realistic P2P simulation environment - protocols previously tested on it have been successfully implemented [17]. In addition the missing random sampling service was provided by the NEWSCAST protocol which is already implemented on PEERSIM. Both SLAC and NEWSCAST are highly scalable (up-to millions of nodes), robust (recovering from noise and the removal of nodes) and completely decentralized (requiring no centralized services).

In the following sections we describe the original tag-based evolutionary algorithm and the derived SLAC P2P algorithm. We then discuss briefly the application task we tested our system with - the Prisoner's Dilemma game followed by an overview of the PEERSIM P2P simulation environment and the NEWSCAST protocol. We then discuss the SLAC implementation details and present the results of experiments performed. Finally we conclude with a discussion of the results and review some of the open issues that need to be addressed in order to progress towards a valuable deployable implementation.

2. Previous Tags Models

The basic algorithm has been adapted from previous (quite different) simulation work using "tags". This work demonstrates a novel method of maintaining high levels of cooperation in environments composed of selfish, adaptive agents. The emphasis of the previous work has been towards understanding biological and social systems [5]. Tags are markings or social cues that are attached to individuals (agents) and are observable by others, often represented in models by a single number, they evolve like any other trait in a given evolutionary model. The key point is that the tags have no direct behavioural implication for the agents that carry them. But through indirect effects, such as the restriction of interaction to those with the same tag value, they can evolve from initially random values into complex ever changing patterns that serve to structure interactions.

The simulated environments in which tags have been applied have generally been very simple with interactions based on pair-wise games with immediate payoffs. Nevertheless, we have attempted to adapt the salient features of such tag systems for application in P2P networks. These features are that agents:

- Restrict interaction to those with whom they share a group defined by tag value
- Selfishly and greedily optimize by preferentially copying the behaviour and tag of others with higher utility
- Periodically mutate their tags and behaviours

By copying and mutating tags, agents effectively move between interaction groups. By restricting interaction within groups free riders tend to kill (reduce the membership) of their own group over time because exploited agents will tend to move elsewhere to get better payoffs, while cooperative groups tend to spread via mutation of the tag. Previous tag models have demonstrated high levels of cooperation in “commons tragedy” [6] scenarios (e.g. in the Prisoners Dilemma – see below). We will not cover the results of the previous tag models in detail here, since the emphasis is not relevant and space precludes detailed treatment, rather we will present our newly derived algorithm (based on the salient features outlined above) and the results we obtained when applying it to two different simulated P2P scenarios.

3. Cooperation and the Prisoner’s Dilemma

Distributed P2P applications often require that nodes behave cooperatively or altruistically to help others in the network. For example, in a file-sharing system, nodes are required to host and upload files on demand to other nodes that require them. Also they need to reply to queries concerning what files they host. But why should nodes do this? In an open system there is an incentive for nodes to behave selfishly - saving their own storage and bandwidth but using other nodes. This problem is not limited to file-sharing because any application that requires other peers to perform actions on their behalf, in some sense, relies on a degree of cooperation from those others. Obviously cooperative behaviours can be built into the peer client software but in an open system how can we ensure that such software will not be changed?

The fundamental issue, then, is: *how can one maintain cooperative (socially beneficial) interactions within an open system under the assumption of high individual (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. Economic and social scientists have often deployed this minimal game as a canonical form of the contradiction that can arise between individual and collective interests.

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

Table 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

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

Prisoner’s Dilemma (PD). When both players cooperate this represents maximizing 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 [10]) captures the intuition that a utility maximizing player would always defect in such games because whatever the other player does a higher score is never attained by choosing to cooperate.

In the context of a P2P system how do we solve this problem without going back to centralized control or closed systems? In the following section we describe the “tag” inspired SLAC algorithm.

4. The SLAC Algorithm

The SLAC algorithm [4] assumes that peer nodes have the freedom to change behaviour (i.e. the way they handle and dispatch requests to and from other nodes) and drop and make links to nodes they know about. In addition, it is assumed nodes have the ability to discover other nodes randomly from the network, compare their performance against other nodes and copy the links and (some of) the behaviours of other nodes.

As discussed previously we assume that nodes will tend to use their abilities to selfishly increase their own utility in a greedy and adaptive way, that is if changing some behaviour or link increases utility then nodes will tend to select it. The algorithm relies on Selfish Link and behaviour Adaptation to produce Cooperation (SLAC) - a task domain independent outline is given below.

Over time nodes engage in some activity and generate some measure of utility U . This might be number of files downloaded or jobs processed etc, depending on the domain. Periodically, each node (i) computes its performance and compares this against another node (j), randomly selected from the population. If $U_i < U_j$ node i drops all current links and copies all node j links and adds a link to j itself - see Figure 1.

Also, periodically, and with low probability, each node adapts its behaviour and links in some randomized way using a kind of “mutation“ operation. Mutation of the links involves removing all existing links and replacing them with a single link to a node randomly drawn from the network. Mutation of the behaviour involves some form of randomized change - the specifics being dictated by the application domain. Previous tag models, on which SLAC is based have indicated that the rate of mutation applied to the links needs to be significantly higher than that applied to the behaviour - by about one order of magnitude [3].

When applied in a suitably large population, over time, the algorithm follows a kind of evolutionary process in which nodes with high utility tend to replace nodes with low

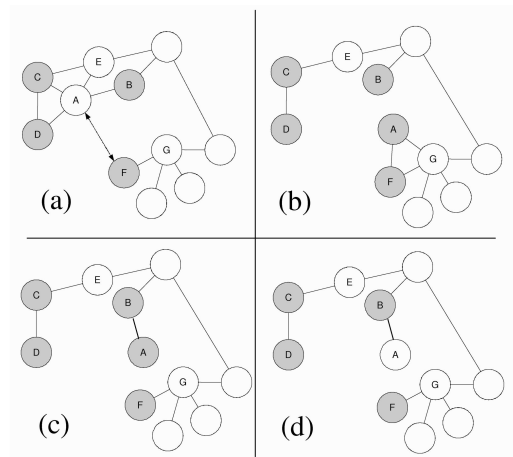


Figure 1. An illustration of ‘replication’ and ‘mutation’ as applied in the Selfish Link-based Adaptation for Cooperation (SLAC) algorithm from [4]. Shading of nodes represents strategy. In (a) the arrowed link represents a comparison of utility between A and F. Assuming F has higher utility then (b) 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).

utility (with nodes periodically changing behaviour and moving in the network). However, as will be seen, this does not lead to the dominance of selfish behaviour – as might be intuitively expected – since a form of incentive mechanism emerges via a kind of ostracism in the network.

5. The Peersim System

Evaluating the performance of P2P protocols is a complex task. One of the main problems for their evaluation, is the extremely large scale that they may reach. P2P networks involving hundred of thousands of peers (or more) are not uncommon (e.g., about 5 millions machines are reported to be connected to the Kazaa/Fasttrack [15] network). In addition P2P systems are highly dynamic environments; they are in a continuous state of flux, with new nodes joining and leaving (or crashing).

These properties are very challenging to deal with. Evaluating a new protocol in a real environment, especially in its early stages of development, is not feasible. Distributed planetary-scale open platforms (e.g., Planet-Lab [16]) to develop and deploy network services are available, but these solutions do not include more than about 500 (at the time of writing) nodes. Thus, for large-scale systems, a scalable simulation test bed is mandatory.

The Peersim P2P simulator [13] has been developed with the aim to deal with the previously stated issues. Its first goals are: extreme scalability and support for dynamism. It is a GPL open-source Java based software project. Peersim has proved to be a valuable tool and it is used as the default experimentation platform in the BISON project [14]. In the following, we provide a brief description of its characteristics.

5.1. Peersim Design Goals

The Peersim simulator is inspired by mainly two objectives:

- High scalability: P2P networks may be composed by millions of nodes. This result can be achieved only with a careful design of the data structures involved, trying to avoid (when possible) any overhead. But the memory footprint is not the only problem: the simulator engine must be also efficient.
- Support for dynamism: the simulator must manage nodes joining and leaving the network at any time; this feature has tightly relations with the engine memory management sub-system.

Another important requirement is the *modular* or *component* inspired architecture. Every entity in the simulation (such as protocols and the environment related objects) must be easily replaceable with similar type entities.

The Peersim extreme performances can be reached only accepting some relaxing assumptions about the simulation details. For example, the overhead introduced by the low level communication protocol stack (e.g., TCP or UDP) is not taken into account because of the huge additional memory and CPU time requirements needed to accomplish this task.

5.2. Peersim Architecture

As previously stated, Peersim is inspired by a modular and very configurable paradigm, trying to limit any unnecessary overhead. The simulator main component is the *Configurator* entity targeted to read configuration files. A configuration file is a plain ASCII text file, basically composed by key-value pairs. The Configurator is the only not interchangeable simulation component. All the other entities can be easily customized.

In a Peersim simulation, the following three distinct kind of elements can be present: protocols, dynamics and observers. Each of them is implemented by a Java class specified in the configuration file. The network in the simulation is represented by a collection of nodes and each node can hold one or more protocols. The communication between node protocols is based on method calls. To provide a specific kind of service, each component must implement a specific *interface*. For example a protocol has to implement at least the `Protocol` or `CDProtocol` interface to run on Peersim.

Peersim has an utility class package to perform statistic computations or to provide some starting topology configuration based on well know models (such as: random-graph, lattice, BA-Graph, ...).

The *Simulator* engine is the component that performs the computation; it has to run the component execution according to the configuration file instructions. At the time of writing, Peersim can perform simulation according to the following execution models:

- Cycle based: at each step, all nodes are selected in a random fashion and each node protocol is invoked in turn;
- Event based: a support for concurrency is provided. A set of events (messages) are scheduled in time and node protocols are run according to the time message delivery order.

This paper work is based on the first simulation model.

```

while(TRUE) do
wait( $\Delta t$ );
neighbour = SELECTPEER();
SENDSTATE(neighbour);
n_state = RECEIVESTATE();
my_state.UPDATE(n_state);

```

(a) Active Thread

```

while(TRUE) do
n_state = RECEIVESTATE();
SENDSTATE(n_state.sender);
my_state.UPDATE(n_state);

```

(b) Passive Thread

Figure 2. The gossip paradigm.

6. the Newscast protocol

Newscast [9] is a gossip-based topology manager protocol. Its aim is to continuously rewire the (logical) connections between hosts. The rewiring process is designed in such a way that the resulting overlay is very close to a random graph. The generated topology is thus very stable and provides robust connectivity. This protocol has been used successfully to implement several P2P protocols, including broadcast [9] and aggregation [8].

As in any large P2P system, a node only knows about a small fixed set of other nodes (due to scalability issues), called *neighbours*. In Newscast, the neighbourhood is represented by a partial, fixed c size view of node *descriptors* composed by a node address and a logical *time-stamp* (e.g., the descriptor creation time).

Referring to the usual gossip scheme (see Figure 2), the protocol behaviour performs the following actions: selects first a neighbour from the local view, exchanges the view with the neighbour, then both participants update their actual view according to the received view. The data actually sent over the network by any Newscast node is represented by the node's own descriptor plus its local view.

In Newscast, the neighbour selection process is performed in a random fashion by the SELECTPEER() method. The UPDATE() method is the Newscast core behaviour. It merges (\cup operation) a received view (sent by a node using SENDSTATE()) with the current peer view in a temporary view list. Finally, Newscast trims this list to obtain the new c size view. The node descriptors discarded are chosen from the most "old" ones, according to the descriptor time-stamp. This approach changes continuously the node descriptors hold in each node view; this implies a continuous rewiring of the graph defined by the set of all node views. This behaviour is shown in Figure 3.

Even though the system is not synchronous, we find it convenient to describe the gossip-scheme execution as a sequence of consecutive real time intervals of length Δ (see the "wait" statement in pseudo-code in Figure 2), called *cycles* or *rounds*.

The protocol always tends to inject new informations in the system and allows an automatic elimination of old node descriptors using the aging approach. This feature is particularly desirable to remove crashed node descriptors and thus to repair the overlay with minor efforts. In addition, the protocol does not require any clock synchronization, but only that the timestamp of node descriptors in each view are mutually consistent.

The topology generated by Newscast has a low diameter and it is close to a random graph having out-degree c . Experimental results proved that a small 20 elements partial view is already sufficient for a very stable and robust connectivity, regardless of the network size.

Newscast is also cheap in terms of network communication. The traffic generated by the protocol involves the exchange a few hundred bytes per cycle for each peer and is estimated in [9].

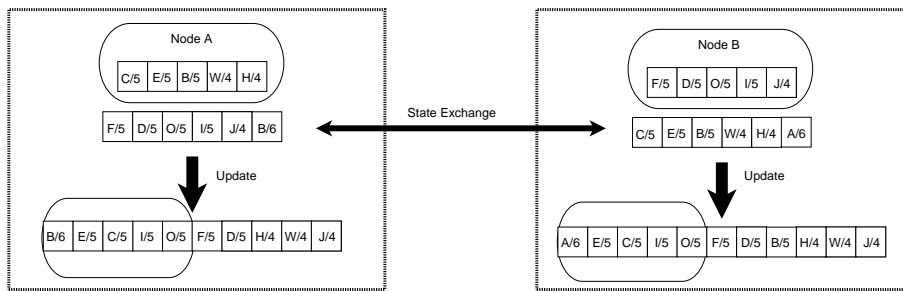


Figure 3. A Newscast exchange between node A (active) and B. Each node has its own 5 descriptor elements view depicted inside the ellipses. A descriptor is a *node-ID, timestamp* pair. After the state exchange node A has received the node B view and viceversa; then each participant merges the received view with its own. The result is depicted under the empty arrow: each node has selected the “freshest” descriptors at random and has discarded the others (those inside the ellipse) to obtain new 5 element view. Note that in this basic example, each node sends its entire view; however, the view can be purged by “old” descriptors before sending.

A protocol such as Newscast provides a service to pick random nodes from the whole network and we can call it *Randomizer Service*. The chance to extract a fresh new node, selected at random from the whole network, is a high desirable source of information for P2P protocols. We can consider such a service as a *building block* for many P2P protocols. In this vision, the cost effectiveness of Newscast is very useful, because the *Randomizer Service* has to be always-on and run by all peers involved in the overlay.

Such a randomizer service can also be a key component during the initialization phase (*bootstrap*) for any higher level protocol in order to fill its view at the beginning. We use Newscast both as a randomizer service and as a bootstrap facility in the Peersim implementation of the SLAC algorithm. We describe this implementation in the next section.

7. Implementing SLAC in Peersim

Now we purpose to implement and test the solutions proposed by the SLAC algorithm into the Peersim platform, on top of Newscast. We have already seen that Newscast has lots of desirable properties such as scalability, robustness and the ability to maintain a random topology. This is why we based our experiments on top of such protocol. Newscast has been previously implemented within Peersim (as previously discussed) and therefore we do not discuss implementation details here.

For SLAC we implemented three classes in Peersim whose names are: **PdDistributionInitializer**, **PdObserver** and of course **PdProtocol**. The first one initializes the nodes of the network with a strategy (cooperate or defect), the second one calculates and prints the results from the simulations. **PdProtocol** is the core of our SLAC implementation and hence deserves a more complete discussion.

The PdProtocol (an overview)

As we have already seen, Peersim is highly modular. So it wasn't so hard to implement a new protocol based on the NetWorld [1] model specifications (which implements SLAC)

compatible with its structure. Peersim simulations [13] are performed through a series of cycles and in each of these some operations are performed. During the simulation our protocol is involved in three phases:

- phase 0: initialization of an auxiliary array for neighbours list (nodes are provided by Newscast);
- phase 1: the PD game is played and the appropriate payoffs are distributed;
- phase 2: two nodes are randomly chosen from the network (these are provided by Newscast) and their payoffs are compared (this is the reproduction phase).

Phase 0 is performed at the first simulation cycle: each node just copies its neighbour list from its Newscast's instance.

In phase 1 the PD game is played between nodes. Each node is initialized with a random value chosen from a set of two (True = Cooperate, False = Defect) and at each cycle each node plays a round of the PD game with one of its randomly chosen neighbours. After this game interaction, payoffs are calculated and distributed.

With phase 2 is performed the reproduction task. It takes place every I cycles (in our experiments we have $I = 4$). In this phase we want to compare the average payoffs of two nodes chosen within the network. The first node we take into account is randomly chosen within the network; the second one, as in phase 0, is taken from Newscast: it is a randomly chosen neighbour of the Newscast instance of the first node. Once two nodes (i, j) are selected this way, reproduction phase can start. Let's consider the case in which node i has a higher average payoff than j , the following operations are performed (obviously if j 's payoffs is greater than i 's the symmetric rules apply):

- all j 's neighbours delete their links to j itself; ¹
- j 's neighbours list is cleared;
- j 's neighbours list is filled in with new items: i 's neighbours;
- the winner itself (i) is added to j 's neighbours list;
- strategy is copied from i to j ;
- some variables held by j are cleared (e.g. number of games played);
- j is added to the respective neighbours lists of the winners neighbours;
- j itself is added to i 's neighbours list;
- mutation is applied with a certain probability.

Neighbours lists have a fixed size F , hence it is not permitted to add nodes when such a limit is reached. That is why this operation is performed by an appropriate function which ensures this limit is not exceeded: before adding a new node, a check on the actual size of the list is performed; if it is equal to F , a randomly chosen node is deleted and then the new one is added, else, if the actual size is smaller than F , the new node is simply added.

This is what happens when i has an average payoff greater than j . Of course if j had an average payoff greater than i , the same algorithm will be performed but with i and j in inverted positions. When the two nodes hold the same value the winner is randomly chosen between the two nodes and then the same operations listed above are performed.

¹Actually, in this implementation of SLAC, this first action was not implemented, interestingly, this change did not stop high cooperation from emerging.

Table 2. Simulation parameter.

Parameter	Value
Cycles	1000
Network size	4000,8000,12000,16000
Initial topology	Random
Maximum degree	20
Strategy mutation rate (MR)	0.001
Tag mutation rate (MRT)	0.001
Reproduction interval (I)	4

Table 3. PD payoffs adopted in the model.

1st player	Strategies	2nd player
1	DC	0
0.8	CC	0.8
0.1	DD	0.1
0	CD	1

At the end of each reproduction step *mutation* is performed. It is applied with low probability (*mutation rate*, $MR = 0.001$), to the node losing the comparison at the beginning of the reproduction phase. Mutation is applied to both strategy and tags (neighbor list). When applied to the strategy bit, it implies the flipping of such bit; when applied to the neighbor list, it involves the cancellation of all the links in the node's neighbors list and the substitution of such nodes with just one node that is randomly selected from the entire network.

8. Simulation Experiments and Results

A series of experiments were done with this new "tag-based" protocol on Newscast. Obtained results relate to experiments done varying several parameters: size of the network, simulation seed, strategy initialization.

The main parameters used are illustrated in Table 2 while the PD payoffs used are those in Table 3.

In the next subsections we will show some of these results and draw some conclusions.

8.1. Cooperation with different network sizes

Results shown in Figures 4 and 5 relates experiments done on networks having different sizes.

The diagram in Figure 4 represents the number of cycles needed to obtain high levels of cooperation (about 93%) over a series of 1000 cycles. The diagram compares the results obtained performing the experiments on four different network size. The results shown in the figure represents the average and standard deviation (out of N sims) of the

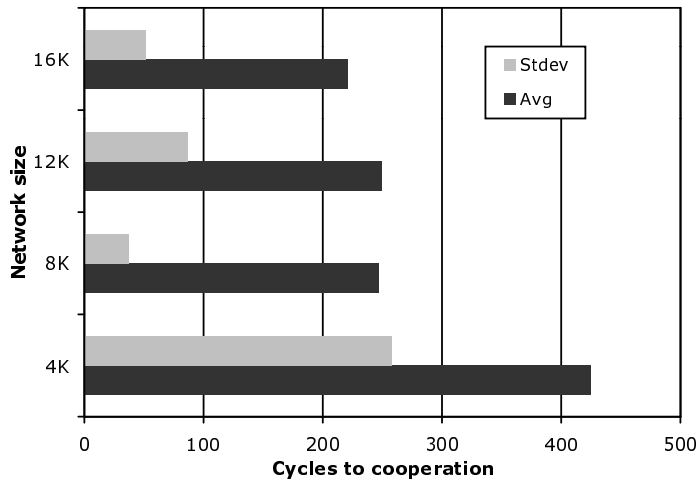


Figure 4. Average number of cycles needed to obtain high levels of cooperation (about 93%) and the relative standard deviation with four different seeds. $MR = MRT = 0.001$, $I = 4$, payoffs from table 3. Results from different network size are compared.

number of cycles needed to obtain cooperation. The network was started from complete defection, the mutation rate used was the same both for the strategy and the neighbour list ($MR = MRT = 0.001$) and the number of cycles occurring between a reproduction phase and the next one was $I = 4$. Let's note the average cycles number for a network of 4000 nodes: it is much higher than the average for the other network sizes and this is because with seed 1, the 93% of cooperation was obtained just after 772 cycles. The same experiment was also performed with a different seeds table but the results we obtained are nearly identical to those just given.

Figure 5 gives the percentage of cooperating nodes over a series of 4000 cycles. The figure just shows the first 150 cycles since after that point there are no relevant changes in the results. Even here the mutation rate was the same both for the strategy and the tag, experiments were performed over different network size and parameter I was set to 4. On the contrary of the previous experiment, here was used always the same seed, hence the percentages we give are not an average. Nodes were initialized at random: at the beginning of the simulation we had a population composed of about half cooperating nodes and half defecting nodes; after the first few cycle, the percentage of cooperating nodes decreased but soon after the 23th cycle it started increasing toward good levels of cooperation.

Observations

From both the diagrams it's easy to note the good level of scalability of the model we are testing: results of Figure 4 are an average of results obtained with different seeds and we succeeded in obtaining similar results for three different network size. The only difference is found with the network composed of 4000 nodes where we obtained a very

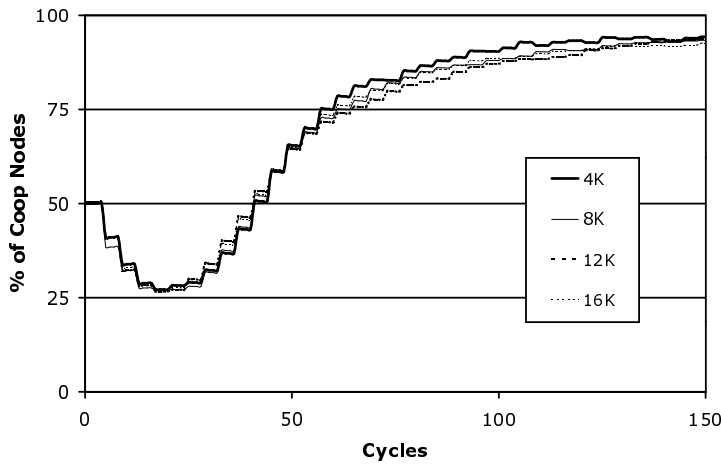


Figure 5. Cooperating nodes over a series of 4000 cycles (single run). $MR=MRT=0.001$, $I=4$, payoffs from table 3. Nodes are initialized with a random strategy. Results after cycle 150 do not increase significantly. Note: the “staircasing” effect is an artefact of the synchronous reproduction at every 4 cycles - with asynchronous reproduction the artefact is not visible.

high value with just one seed. It would be interesting to make further experiments with more seeds. Figure 5 gives the same important result: here using always the same seed, we obtained the same trend with all the sizes.

8.2. Cooperation with long runs

Some experiments were done with a big number of cycles. On a network composed of 4000 nodes were performed a series of experiments; for each of them we used a different random seed, $MR = MRT = 0.001$, $I = 4$ and performed 10000 cycles.

Results are very close for each seed used, hence in Figure 6 we just propose those obtained with one of them. The diagram shows that good levels of cooperation can be obtained from cycle 645 (95.5 %): from this cycle to cycle 10000 the average percentage of cooperating nodes is 95.40 and the standard deviation is 1.03.

Observations

The experiments just proposed have a great importance since they test the reliability of the model in the time. From all the tests we made (some even with 50000 generations), and also in those proposed in this section, we found that once cooperation has started, it never claps and can be sustained for long times; we also learned that cooperation can be sustained at good levels.

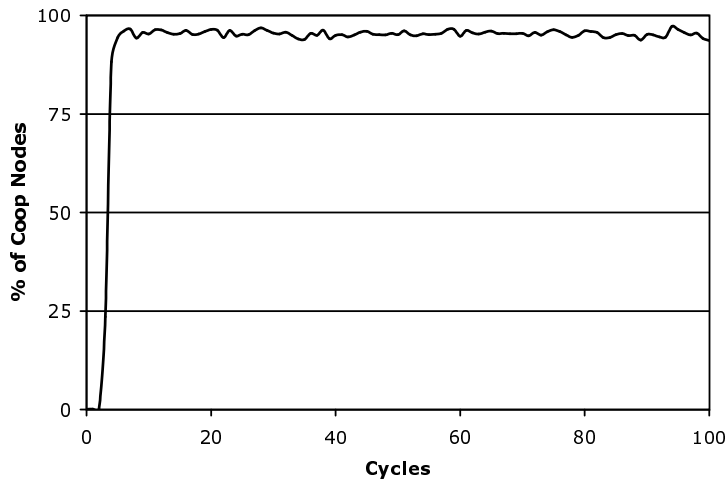


Figure 6. Cooperation with a series of 10000 cycles (single run). Nodes = 4000, MR=MRT=0.001, I=4, pay-offs from table 3.

9. Discussion and Conclusion

The results we obtained indicate that high cooperation is produced when nodes follow the SLAC algorithm. Even though the SLAC algorithm implements nodes that behave selfishly in a myopic and greedy way - that is, they copy other nodes in the network that have higher utility - high levels of cooperation are produced in the single round Prisoner's Dilemma (PD) game.

These new results therefore confirm those results previously found [2,4] in similar simulation experiments and this adds confidence that SLAC is robust to different simulation implementation details. Here, for example, we interleaved the reproduction phase with the interaction phase whereas in previous simulations reproduction followed at the end of each cycle of interaction. Also in further experiments (not shown here) we found that similar results were obtained when reproduction was interleaved with interaction in a fully asynchronous way - where each node has a probability of reproduction after interaction.

Interestingly, the results shown in Figure 4 appear to recover some of the reverse scaling properties demonstrated in an early non-network based tag model [5] which appeared to have been lost in our initial network model [2]. However, more runs and analysis are needed to explore this question. However, we can certainly state that in all the experiments so far performed with the SLAC algorithm larger networks do not take longer to converge and often converge more quickly. This is obviously a valuable property for any candidate algorithm for large scale systems.

A further finding of these new results appears to contradict earlier generalizations [3] that tag-type models needed to have higher mutation rates on the "tag" than the strategy - in the case of SLAC this would mean a higher mutation rate on the neighbour list or view which contains the links to neighbour nodes than the behavioural strategy of the node (either to cooperate or defect). But here high levels of cooperation were produced when

the mutation rate was the same. This indicates that further work is needed to circumscribe such a generalization since it is currently unclear what difference in implementation has allowed this assumption to be relaxed.

These results also demonstrate that the NEWSCAST protocol can be used to provide the random sampling service required by SLAC but not previously explicitly modelled in simulation. This is important since any actual implementation of SLAC must have access to such a service that is both scalable and robust. NEWSCAST provides such a service[9] with the additional benefit that it has actually been tested in the form of a real implementation [17].

We have argued, and demonstrated previously, that cooperation in the single round PD indicates that cooperation can be produced in other more realistic task domains [4]. We are therefore confident that these results indicate that the PEERSIM implementation could support cooperation in other task domains (such as file sharing or other kinds of resource sharing).

Finally, we note two major issues that could destroy cooperation within SLAC. Firstly, we currently assume that nodes are able to compare utilities correctly, that is, we assume nodes report their utilities honestly when requested to do so by nodes. But what would happen if nodes lied about their utilities or just failed to report anything? This introduces a kind of "second order" free-rider problem at the informational level because if we assume nodes may behave selfishly and / or maliciously then we need to demonstrate individual incentives for supplying correct utility values. Secondly, we also assume nodes will allow themselves to be copied by supplying their behaviour strategy and their current neighbour list or view (containing their node links) to other nodes. Again, this may not be case with malicious and selfish nodes in certain contexts. Both of these issues we aim to address in future work.

Acknowledgements

This work would not have been possible without perceptive discussions with many people, particularly those in the Bologna group including: Mark Jelasity, Alberto Montresor and Simon Patarin. Additionally, we thank the anonymous reviewers of the initial draft of this paper for their comments and suggestions.

References

- [1] D. Hales, B. Edmonds Applying a socially-inspired technique (tags) to improve cooperation in P2P Networks *IEEE Transactions in Systems, Man and Cybernetics - Part A: Systems and Humans* pp.385-395 2005
- [2] D. Hales. Self-Organizing, Open and Cooperative P2P Societies From Tags to Networks. *Proceedings of the 2nd Workshop on Engineering Self-Organizing Applications (ESOA 2004)*, LNCS 3464, pp.123-137. Springer, 2005.
- [3] D. Hales. Change Your Tags Fast! – a necessary condition for cooperation? *Proceedings of the Workshop on Multi-Agents and Multi-Agent-Based Simulation (MABS 2004)*, LNAI 3415. Springer, 2005.
- [4] D. Hales. 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. Available at: <http://www.davidhales.com>
- [5] D. Hales. Cooperation without Space or Memory: Tags, Groups and the Prisoner's Dilemma. In *Moss and Davidsson (eds.) Multi-Agent-Based Simulation*. LNAI 1979:157-166. Springer. Berlin. 2000
 - [6] G. Hardin. The tragedy of the commons. *Science*, 162, 1243-1248. 1968
 - [7] J. Holland. The Effect of Lables (Tags) on Social Interactions. *Santa Fe Institute Working Paper 93-10-064*. Santa Fe, NM 1993
 - [8] M. Jelasity and A. Montresor, Epidemic-Style Proactive Aggregation in Large Overlay Networks, in *Proceedings of the 24th International Conference on Distributed Computing Systems (ICDCS'04)*, March 2004, pp. 102–109, IEEE Computer Society, Available at: <http://www.cs.unibo.it/bison/publications/icdcs04.pdf>
 - [9] M. Jelasity and W. Kowalczyk and M. van Steen Newscast Computing *Technical Report IR-CS-006*, Vrije Universiteit Amsterdam, Department of Computer Science, November 2003, Available at: <http://www.cs.vu.nl/globe/techreps.html#IR-CS-006.03>
 - [10] J. F. Nash. Equilibrium Points in N-Person Games, *Proc. Natl. Acad. Sci. USA* 36, 48-49, (1950).
 - [11] R. Riolo, M. D. Cohen, R. Axelrod. Cooperation without Reciprocity. *Nature* 414, 441-443, (2001).
 - [12] R. Tivers. The evolution of reciprocal altruism. *Q. Rev. Biol.* 46, 35-57. 1971
 - [13] Peersim Peer-to-Peer Simulator, Available at: <http://peersim.sf.net>
 - [14] The BISON Project, <http://www.cs.unibo.it/bison>
 - [15] Kazaa Web Site, <http://www.kazaa.com>
 - [16] PlanetLab Planetary-Scale Testbed, <http://www.planet-lab.org>
 - [17] T. Binci. EpidEm: EPIDemic EMulator, Graduate Thesis in Computer Science, University of Bologna, Department of Computer Science. Available at: <http://bincit.web.cs.unibo.it/index.htm>