# KNOWLEDGE INTEGRATION AND FREE RIDING IN LARGE ORGANISATIONS: EXPLORING GOVERNANCE MECHANISMS, BEHVAIOUR AND PERFORMANCES

xxxxxxx xxxxxxx and xxxx xxxxx xxxx
Dept. of xxxxxxxx xxxxxxx,
University xx xxxxxxx, xxxxx
email: {xxx , xxx}@xx.xxxxx.xx

January, 2008

**Abstract**

A prevalent claim is that we are in a *knowledge economy.* An increasingly influential argument is that the division of labour is becoming complex and firms can be viewed as networks of knowledge nodes, that is, sets of interacting individuals with key skills and competencies. Firms competitive advantage relies in their ability to enact intellectual production processes that require bridging talents within an organisation. Building intra-organisational networks of skills, however, may request the enactment of social processes of leading to the sharing of competencies and the exchange of information. The issue, then, becomes one of understanding the circumstances under which firms provide an appropriate context to promote cooperation among individuals bearing specific skills and to mould social networking within organisations. By the means of simulation experiments, we investigate how cooperation may emerge in organisational contexts. In particular, we focus on the role of imitation of pro-social policies and explore role played by scope of interaction dynamics in fostering cooperation.

## 1 Introduction

A prevalent claim is that we are in a *knowledge economy.* In this work, we take the view that what characterises a knowledge economy is the growing importance of human capital in productive processes [9] and the increasing knowledge intensity of jobs [16]. In addition, an increasingly influential argument is that the division of labour is becoming complex and firms can be viewed as networks of knowledge nodes [9], that is, sets of interacting individuals with key skills and competencies where firms competitive advantage relies in their ability to enact intellectual production processes that require sharing competencies, bridging talents within an organisation and exchanging information. Such networks crystallises firm-specific knowledge and provide ground upon which firms build their heterogeneity. Thus, the fact that the knowledge content of jobs increases raises questions concerning emerging organisational forms. In particular, assuming that economic actors are boundedly rational and stocks of knowledge of two individuals are different, cognitive limitations exist so that the skills of one person cannot be easily absorbed by or transferred to another person (see [5, 11]). Second, intellectual production processes often do not only require simple exchange of discrete and separable products, but they require joint production between individuals bearing specific skills ([5]). Thus, performing a specific task may entail coordination of simultaneous efforts and the enactment of complex social processes of interaction among specialists bearing specific skills. In intellectual joint production processes, however, especially when number of individuals involved increases, incentive to shirking may emerge. Traditionally, firms are considered as an appropriate mechanism to control opportunism and induce coordination of collective action by the means of hierarchy. As, however, production processes become increasingly intertwined with social processes, a doubt may emerge regarding the ability of firms to use hierarchy to influence such social processes. In this paper, we explore how organizational context, intra-organisational networks of specialists and individual decision-making interact to produce emergent cooperation. The issue, then, becomes one of capturing the means that hierarchy within firms has to mould social networking within organisations. As Adler suggests ([2]) as knowledge becomes more important in our economies, increasingly transactions will be managed relying on trust-based organizational forms. In this respect, we need to explore whether a firm, rather than other organisational form, is more able to produce cooperation and coordination among agents. That is, given the weakness of top down hierarchical control, to what extent a firm, better than other organizational modes, create trust and cooperation thereby controlling agents opportunism? We focus on the problem of cooperation within organizations. In particular, we address the social dilemmas emerging when individuals have a choice between adopting a prosocial and a selfish course of action within an organization. We build an agent-based model to simulate interaction among individual strategies, network structure in which agents are embedded and organizational context. We characterize two types of agents. First, we define a *prosocial* agent who does not withdraw its effort and pursues its utility by searching coordination. Second, we define an agent, which we call *free raider*, who maximizes its utility by withdrawing its effort while yielding same reward. By the means of simulation experiments, we investigate how cooperation may emerge in organisational contexts

# 2 Transaction costs economics. Opportunism and the firm

Transaction costs economics suggests that firms exist to control opportunism by the use of hierarchy. This approach is very useful to explain emergence of firms, and boundaries among firms, when means of production are non-human. Non-human means of production, once owned, can be directed to perform specific tasks. For example, if a supplier has bargaining power because she owns a physical resource available in limited supply, by acquiring the supplier a firm can control opportunism of the supplier and have direct control on the physical resource and on the productive process that this latter underpins. Hierarchical control is less effective in presence of means of productions crystallised in human beings especially in a situation of small numbers, that is, a situation in which a firm needs to control a professional with unique and valuable skills. In this case, it is unlikely that hiring the professional entails effective control. He can abandon the firm as he wants and he will always find another firm eager to hire him. Share in the equity or stock options are remedies that are only partially effective to control opportunism as agency literature suggests. This is a classical problem of collective action; as Olson suggests, the fact that rational actors have common goal (in this case the value produced within a firm) does not necessarily implies that the actors work at their best to achieve the goal ([21]). In general, agents endowed with bargaining power arduously can be controlled solely with hierarchy. Contracts can be signed that link reward to performances but if the agent has enough bargaining power, and is enough opportunist, these contracts may ineffective in forcing the agent to provide its maximum effort. For example, the contracts may contain objectives easily achievable with limited effort. In this case a typical solution is to internally develop unique valuable resources ([19]). Yet, for unique competencies that are characterised by large components of tacit and uncodified knowledge, reproducing the valuable skill may be a problem. In addition, when productive processes entail social interaction, it may be difficult to reproduce a successful experience by reconstructing a subtle network of personal relationship and group commitment.

Assuming that supply of talented employees is not limited, and the small number problem does not apply, still using hierarchy to control opportunism in employment relationship may be difficult as well. Difficulty in controlling transactions increases with two factors. First, it may become difficult to assess whether the execution of the task is appropriate. Two interacting elements may increase difficulty in evaluating task execution. First element is the ambiguity of the content of a task, and second element is the specialisation of skills employed in a productive process. As the content of a task becomes unstructured, the content of a task becomes more ambiguous and the contract is increasingly incomplete. In addition, as specialization of skills increases, a *supervisory problem* (see [16]) is pending because nominal supervisors will not know the best way of doing the job - or even the precise purpose of the specialist job itself - and the worker will know better. Thus, by internalizing within an organization task execution firm negotiate with the employees the possibility, in exchange of a salary, to use hierarchy to direct employees behaviour by selecting an action over a range of possible actions, these latter referred to as the area of acceptance. Yet, as the contract becomes more incomplete, the boundaries of the area of acceptance becomes fuzzy and skills become highly specialized, it may be difficult to assess whether an employees behaviour reflect the content, implicit or explicit, of a contract or whether the behaviour applied to the task is the optimal one or the one that minimize employees effort. Second, as intellectual production processes and such intellectual production processes require sharing competencies, bridging talents within an organisation and exchanging information, production requires the enactment of social processes of networking. Thus, increasingly, within organizations, along with formal governance mechanisms, dynamics of intra-organizational social networks are responsible for the performance of intellectual productive processes. In general, intellectual production processes can be considered largely as peer-production processes in which individuals own large portions of the (intellectual) means of production employed and can arduously be directed by the means of hierarchy. In the peer-production, exchanges are less negotiated and more grounding upon reciprocity. In addition, evolving ICT technology allows dynamic group formation, aggregation in virtual teams and interaction through intra-organizational and extra-organization virtual environments.

## 3 The issue of emerging cooperation

The analysis of situations in which individual rational choices lead to globally undesired outcomes has been at the core of a large body of literature dealing with social dilemmas ([21, 15]). The issue of cooperation has been of interest for economists, game theorists, sociologists, anthropologists and biologists in the last few decades. The appeal of the topic for economists is connected with the following puzzle: why do selfish individuals cooperate rather than simply maximizing their utility? Interestingly, anthropology, on one side, and biology, on the other side, provided theories and concepts to explain the emergence of cooperation in different ways. On the one hand, anthropologists have focused on the concept of *reciprocity* and the inclination of human beings to reciprocate gifts ([25, 10, 22]). On the other hand, biologists proposed theories suggesting that altruism is long-term material self-interest. The theory of *inclusive fitness* ([13]) explains cooperation as the result of altruistic behavior towards relatives; in other words, what is apparently an altruistic behaviour is an adaptive mechanism that preserves ones kinship. The theory of *reciprocal altruism* ([30]) proposes that individuals behave altruistically expecting reciprocation. This ability to contemporaneously explain cooperative behaviour and to preserve the hypothesis of individual behavior as fundamentally selfish is what intrigued economists who adopted these biological theories, and especially the theory of reciprocal altruism, to explain evidences of cooperation. In this respect, cooperation emerges as the result of a repeated interaction among two players who are bound to interact within the framework of a Prisoner Dilemma and have the possibility to punish each others defection ([29]). Along these lines, starting from the contribution of Axelrod ([3]), emerging cooperation has been investigated by simulating repeated, multi-person Prisoner Dilemma game ([23, 26, 28, 20, 27]). Within the context of a repeated, multi-person Prisoner Dilemma, the role played by structural embeddedness in explaining emergent cooperation among individuals has been highlighted by Macy and Skvoretz ([24]). Only recently, however, the hypothesis that individuals decision-making and social structure coevolve has been taken into consideration. Eguiluz et al ([7]), for example, assumed explicitly *social plasticity*, that is the ability of an individual agent to select partners thereby changing their neighbourhood as time goes by. In their model, player $i$, when imitating a defector, replaces with probability $p$ the link with the imitated defector and another link pointing to a randomly chosen partner from the whole network. In Hanaki et al. ([14]), the interaction dynamics between an individual and the network structure in which it is embedded includes the hypothesis of mutual consent, that is, for a new tie to be created it is necessary that the individual at which the new link is pointing agrees to the creation of the link.

Despite the fact that studies on emerging cooperation have reached a stage of development in which the social structure in which players are embedded receives an explicit role, the existent body of literature does not explicitly deal with the issue of cooperation when social structure emerges within a firm. The focus on firms implies consequences in the approach needed to model emerging cooperation. One key issue we emphasise concerns the mechanism by which individuals imitate others and consequently learn to cooperate. Most of modelling on emerging cooperation includes individual learning processes by which agents imitate better performing strategies. Within firms such possibility to observe others performances is biased by a number of factors. Individuals observe others salaries but salaries change according to organisational mechanisms so that it may be difficult to use others salaries to select candidate strategies to be imitated. For example, salaries may depend on roles and different career ladders, individuals need to select for imitation other individuals that are in comparable career path. In addition, imitating an individual who is earning an higher salary but is producing a much larger effort may be suboptimal. Therefore, imitation and learning implies the possibility to detect and interpret the amount of effort produced within a production process. As a consequence, results of imitation and learning depend on the accuratedness of information available on individuals to be imitated. The prisoner dilemma may be a limited framework to represent interaction dynamics within organizations in which assessment of others behaviour implies an accurate analysis of both salary and effort produced. The main contribution of this paper is to extend the standard modelling framework to include the analysis of the interaction among individual choice, social network and a firms organizational context.

# 4 Methodology

Modelling and simulation constitute a fundamental element of the research design. Simulation helps rigorously to deduce consequences from modelled assumptions when complexity of modelling makes difficult to obtain closed-form solutions. In addition, simulation allows looking at unfolding organisational and social processes, capturing the behavioural characteristics in transitory states. In this work, we use a computer simulation model as a theoretical laboratory to analyse the circumstances in which different hiring and reward strategies, firms heterogeneity and rent distribution patterns emerge. Alternative hypothetical, though dormant, trajectories will be activated by modifying the underlying modelled assumptions. This approach has the advantage of creating an appropriate setting to conduct controlled experiments. History can be re-run, showing how small, *ab-initio* modifications in parameter values can be amplified over time, to yield firms with distinct characteristics. Simulation is a unique methodology to perform this journey in history. This kind of method is a form of computational thought experiment: in which we ask what if questions in an artificial world. However, the ultimate aim is to allow us to develop hypotheses and theories that can then applied to real world phenomena and data. Our ultimate aim is to understand the real world. We use the computer model at this stage to help us to generate and test, in a rigorous and deductive way, candidate ideas.

We built and used an agent-based model the CoopNet model - to simulate interaction among employees. The processing of using computer simulation models in this way (see [4]) is an emerging paradigm within the social sciences. Increasingly social scientists are using the techniques of multi-agent based simulation (MABS) to explore complex dynamics in artificial social systems [12].

The CoopNet model should be viewed as an artificial society type model (i.e., similar to the SugarScape model [8]). The CoopNet model allows use to express formally (computationally) a number of hypotheses about potential processes that may occur in real organisations but in a stylized and executable manner such that experiments can be performed to deduce the consequences of those hypotheses when they are combined in complex, adaptive systems (CAS). We therefore purposefully present a simplified model in which we hope to capture the kinds of complex dynamics in which we are interested.

# 5 Model

The basic unit in our model represents an *hour of work* and we assume a working week consisting of 40 working hours. Fractional units are not allowed.

Each node in the system holds a list of links to other nodes, which we call cache. Therefore, a network defined by the relation "who knows whom" is induced by the cache's information. In addition to the network (or overlay) induced by the CoopNet model, we assume the presence of an underlying random network to ensure a robust connectivity layer in case of partitioning of the CoopNet overlay. Essentially, a node has access to two distinct caches: the random one and the CoopNet one.

We can define more formally both networks as, respectively, a graph $G_{coop}(V, E_{coop})$ and $G_{rnd}(V, E_{rnd})$, where $V$ is the same set of vertices (nodes), while $E_{coop}$ and $E_{rnd}$ are the corresponding sets of links. The graph $G_{coop}$ is undirected, while $G_{rnd}$ is directed.

Each node $i$ holds a variable $p_i$ $(effort)$ representing its current *productivity* effort and a variable $w_i$ representing its *wage*. With $p_i(j)$ we indicate the portion of the effort spent by $i$ in reciprocating with $j$. Each node $i$ can act in a *pro-social* or *free raider* way; essentially, pro-social nodes will try to invest more, while free raider nodes prefer to maximize their utility by withdrawing their effort while yielding the same reward.

Every node, holds a BlackList structure keeping track of those nodes with which he believes or suspects the cooperation is not good enough. To avoid wrong suspicions, the entries in the list are purged on a temporal basis; if the entry is already present when a new suspicion is raised, then the entry's *time-to-live* is reinforced.

In addition, each node is aware of a global variable, *PM*, measuring the average system productivity level: $\sum_{i=1}^{N} \frac{p_i}{N}$, where $p_i$ is the *effort* played by each node and $N$ is the size of the network. However, when the network splits in distinct connected components, $PM$ is computed on a per-component manner. In other words, it is computed on the set of efforts of each component: $\sum_{i=1}^{N_k} \frac{p_i}{N_k}$, where $N_k$ is the size of the k-th connected component. For simplicity, we consider $PM_i$ as the $PM$ of component-i.

The value of $PM$, as well as each node's wage, is resampled by an oracle at regular intervals (i.e., at the end of the play phase, see 6.1). The wage is actually computed as a function of the node's perceived $PM$ following the rules shown in Table 1.

We consider SYS_MAX_EFFORT=50 (hours) as the maximum threshold value

| Condition | Rule | Comment |
|---|---|---|
| $PM > 40$ | $w_i+ = 1.1$ | extra wage |
| $35 \leq PM \leq 40$ | $w_i+ = 1$ | std. wage |
| $20 \leq PM < 35$ | $w_i+ = 0.9$ | light punishment |
| $PM < 20$ | $w_i+ = 0.7$ | punishment |

Table 1: Reward scheme.

for $p$.

In our model, the $G_{rnd}$ graph is built and managed by a specific service which we call *Peer Sampling Service* (PSS) [17]. We just provide here the basic idea of the service of which are available several implementations (see [18, 31]). Essentially, the PSS can build and maintain a robust, random overlay network using a *gossip-based* communication model [6]. In addition, it works in a fully decentralized manner and provides each node with a random sample of neighbor links (hold in each node's CACHE) which is continuously updated over time.

The $G_{coop}$ graph instead is built and managed by the service provided by the CoopNet algorithm discussed in the following section.

# 6 CoopNet algorithm

The CoopNet schema distinguishes among three distinct phases: (a) *play*, (b) *drop* and (c) *rewire*. The simulation scheme lets every node to execute first the *play* phase, then all nodes execute the *drop* phase and so on. An optional 4-th phase, called *firing*, is also present (see Section 6.4).

## 6.1 Play phase

In this phase each node try to obtain the highest possible wage by interacting with its neighbors. If the interaction with one or more neighbor is not satisfactory, then those neighbors can be blacklisted and marked in order to be drop in the next phase. The play phase is iterated for a number of steps $t$; each step represents a

**Pro-social behavior:**

```
∀ neighbor j ∈ CACHE do
  if j ∈ BLACKLIST → j : to be drop
  else if p_i(j) < p_j(i) → CLOSEGAP(i, j)
  else if p_i(j) = p_j(i) → ∅
  else
   CLOSEGAP(j, i)
   wait[j]++
   if wait[j] ≥ MAX_WAIT do
    {j} ∪ BLACKLIST
    j : to be drop
    wait[j] = 0
```

**Free raider behavior:**

```
∀ neighbor j ∈ CACHE do
  if j ∈ BLACKLIST → j : to be drop
  else if p_i(j) > p_j(i) → CLOSEGAP(j, i)
  else if p_i(j) ≤ p_j(i) → ∅
```

Figure 1: The CoopNet behaviors. A behavior code is executed $t$ times during the play phase by each node.

week of work and we typically adopt $t = 4$ to model a month of work.

The perfect satisfaction in a relation is when both parties are spending the same amount of effort. However, the *willingness to cooperate* is a very important factor in our model and therefore we consider satisfactory a relation in which the weakest party shows its will to cooperate (i.e., *partially* fills the effort gap) during a $t$ steps period.

Each node can play using an *cooperative* or *defeating* behavior; each behavior is depicted in pseudo-code in Figure 1. The node's behavior is assigned at bootstrap and does not change over time, if not stated otherwise.

### 6.1.1 Effort distribution and adjust

In the first step ($t_1$) each node $i$ will profuse an effort $p_i(j) = \frac{p_i}{degree_i}$ between each neighbor $j$.

In the next steps , each node have to adjust its effort according to the other party's efforts. The basic idea is that when two nodes $i$ and $j$ plays distinct values, say $x_i$ and $x_j$, they tend to close the difference gap (i.e., the CloseGap function shown in the pseudo code). Closing the gap is a relative concept and it is tightly related to the actual node's behavior. For the ease of the presentation, we present the $\textsc{CloseGap}(i, j)$ function as a monolithic procedure. However, it involves a communication among nodes $i$ and $j$ as any interaction among agents involves. In other words, calling the $\textsc{CloseGap}(i, j)$ function means that *both* nodes $i$ and $j$ have to adjust their effort according to their behavior.

A cooperative node $i$, for instance, will try to play the same amount of a neighbor $j$, but if $j$ has played less than $i$, it will not modify its effort waiting and hoping that neighbor $j$ is willing to play more in the next iterations. Node $i$ will wait node $j$ for a limited amount of time (i.e., `MAX_WAIT` system wide parameter); if node $j$ is cooperative it will likely fill (at least partially) the effort gap.

Conversely, a free raider node will not modify its effort if its neighbor has played more, otherwise it will play less if its neighbor's effort is less than its own effort.

Each neighbor pair has $t$ interaction steps at its disposal to adjust the effort gap among them.

## 6.2 Drop phase

After the *play* phase, each node $i$ checks its neighborhood and simply unlinks those nodes which have been previously marked to be drop. Note that the unlink process is bi-directional.

When a relation is drop, both parties reinvest the amount of effort, previously spent with the old neighbor, among the actual neighbors left. This effort is reinvested in a uniform fashion among the neighborhood.

## 6.3 Rewire phase

As in the *play* phase, the rewiring process is related with the node's behavior. We suppose that node $i$, during the drop phase, has marked a set of size $k$ (where, of course, $k \leq degree_i$) of its neighbors as "to be drop"; the rewiring strategy followed by node $i$ comprises three distinct primitive actions: (1) SEARCH, (2) SELECT and (3) ACCEPT.

The strategy of each action differs depending on the actual node's behavior:

- SEARCH:

  C: pick a random node from the underlying network

  D:
  - * if $degree_i = 0$, pick a random node
  - * otherwise, pick $k$ random nodes from the underlying network if $\exists\, w$ nodes $\in$ CACHE with $k = w$ and $p_i(w) > p_w(i)$, such that lowering the effort of $i$ would lead to $PM_i(t) > PM_i(t+1)$

- SELECT:

  C, D: `<condition>` and $j \notin$ CACHE, where $j$ is the candidate neighbor

- ACCEPT:

  C: if node $i$ has room in its cache; otherwise, if the newcomer $j$ has $p_j > min_i\{p_1, p_2, \ldots, p_{degree(i)}\}$, then $i$ substitutes its worst performing neighbor with j

  D: if node $i$ has room in its cache and $PM_i < 15$

The SELECT strategy can be implemented in several ways using distinct `<condition>` statements. All available choices are summarized in Table 2.

The first two actions define how to search for new neighbor candidates (e.s., randomly) and the condition with which to accept or deny them.

The latter action instead, reflects the P2P nature of the interaction among the parties; each node $i$ has a passive part that is listening for queries (i.e., "I'd like you to be my neighbor" message) from other nodes. Therefore, the ACCEPT action defines when a node $i$ can accept a querying node $j$ as a new a neighbor.

| **Pro-social** | **Free raider** |
|---|---|
| $\mathbf{p_j \geq PM_i}$ | $\mathbf{w_j \geq w_i}$ |
| $\mathbf{p_j \geq PM_i}$ | $\mathbf{p_j \geq PM_i}$ |
| $\mathbf{p_j \geq PM_i}$ | $\mathbf{p_j \geq p_j}$ |
| $p_j \geq p_j$ | $p_j \geq PM_i$ |
| $\mathbf{p_j \geq p_j}$ | $\mathbf{p_j \geq p_j}$ |
| $p_j \geq p_j$ | $w_j \geq w_i$ |
| $\mathbf{w_j \geq w_i}$ | $\mathbf{w_j \geq w_i}$ |
| $w_j \geq w_i$ | $p_j \geq p_j$ |
| $w_j \geq w_i$ | $p_j \geq PM_i$ |

Table 2: Available choices for the <`condition`> statement in SELECT strategy. The rows in bold highlight our preferred combinations.

In ACCEPT, when a node is free raider, it accepts queries only when it is in a non-acceptable condition - i.e., when its effort is too low and he run the risk to be fired -, otherwise he do not accept any join request.

The idea behind this behavior is that a free raider nodes tend to be lazy and to limit their relations with other nodes. Essentially, a free raider node is willing to have just a few good links in order to minimize its effort and to have a satisfying wage as well. In an extreme case it may have no links as long as it is satisfied by its effort and the effort is not low enough to trigger the firing process.

It is important to note that the rewiring process may fail to choose new neighbors, because of the requirements of action (2) and (3) and a node could not be able to renew its drop neighbors. In other words, both parties have to agree to startup a new link (collaboration) according to their strict constraints and the set of possible candidates to choose from is limited by the underlying network view.

### 6.3.1 Effort redistribution during rewire

When a node executes the rewire phase, no matter if actively or passively, he has to grant some effort to the neighbor with which he joins. If the node is a pro-social, he will grant a random effort in the range [0, SYS_MAX_EFFORT]; otherwise, he will grant a random effort in the range [0, $p_i(t_0)$ - $p_i$], where $p_i(t_0)$ is the node's initial effort. In other words, this is the only case in which a free raider node can increase its current effort $p_i$; however, it can never exceed its $p_i(t_0)$ value.

### 6.4 Firing

We assume that the system can fire any node $i$ that is putting an effort $p_i < \delta$. In addition, we consider two distinct firing policies:

- **fuzzy**: the probability to fire a lazy node is proportional to $p$
- **strict**: the firing process is simply deterministic: any node $i$ having $p_i < \delta$ is purged from the system

The threshold $\delta$ is a system parameter, thus it can be fine tuned.

These lazy nodes are substituted by an equal amount of "fresh" nodes whose initial effort $p_i$ is initialized using the same normal distribution used at the system bootstrap.

## 7 Experiments

The experiments described in this section have been performed using a prototype implementation of the model, written in the Java language and running on top of PeerSim [1], an open source P2P simulator. We focus on a scenario in which 500 agents are involved. The setup of each node and its environment is discussed as follows.

The goal of our experiments was threefold: first of all, we focused on the structural characteristics of the emergent CoopNet network; second, we measured the performance of CoopNet in terms of distinct metrics that will be introduced in the followings. Third, we checked if and how some topological properties are related to the performance achieved. Finally, we tested an extended version of CoopNet in which we let each node to modify its behavior according to a specific policy (see Section 7.1).

In our experiments, each agent participates in two distinct overlays: a random overlay, providing robust connectivity and the CoopNet overlay reflecting the work relationships. The links of both overlays are randomly assigned at the bootstrap. The degrees of the networks are respectively set to: $k_{rnd} = 20$ and $k_{cnet} = 5$.

The maximum effort a node can spend in a *time unit* is fixed to the constant SYS_MAX_EFFORT = 50 and it is considered a system wide parameter. For simplicity, we define that the time granularity in our system - which we call *cycle* - corresponds to a week of wall clock time. The effort with which each node starts to play is assigned by a normal distribution having parameters $\mu = 27.5$ and $\sigma = 5.5$. These values are motivated by the fact that we want to start having the vast majority of the population with an average value and few nodes at the extremes.

The node's behavior (i.e., pro-social or free raider) is assigned according to a 50-50% proportion and the default rewiring policy is $ppm - ww$, if not stated otherwise.

When the firing feature (see Section 6.4) is enabled, any new node injected in the system is initialized with the same effort and behavior distribution as described above.

Figure 2 shows some properties of the emerged clusters. On the left picture is shown the number of clusters, while on the right picture the biggest cluster size is depicted. In both pictures the results are presented according to the absence / presence of the firing component.

In the left picture, the network first splits in many distinct clusters because of the diversity of the efforts, in both cases (nofiring/firing scenarios). Then, the number of clusters explodes after cycle 100 and the network lets emerge many

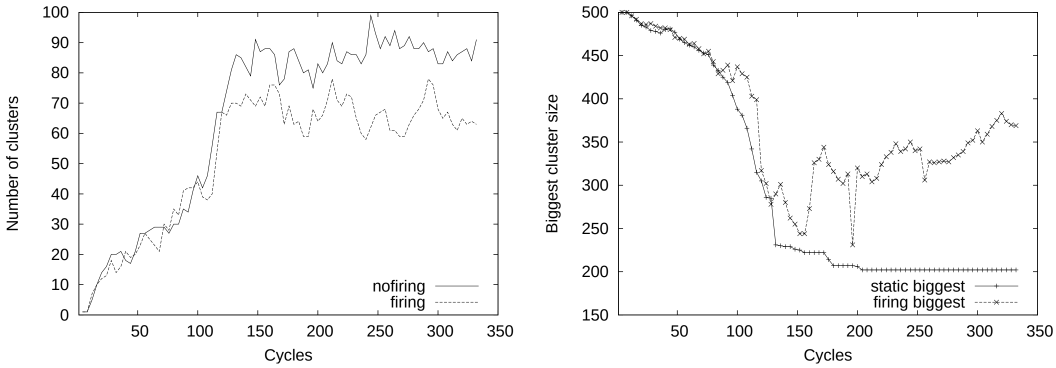


Figure 2: The left picture shows the number of emerging clusters, while the picture on the right depicts the biggest cluster size; both figures show the result according to the presence or absence of the firing component. Network size is 500.

subgroups of nodes having similar (effort) characteristics. As poor performing nodes are eliminated in the firing scenario, the effort is more uniformly distributed and nodes tend to work together more easily without originating other groups (clusters).

The right picture is particularly important as it depicts the formation of a giant component in both cases. In general, the presence of a giant connected component is an interesting property for us, as it means that the relation of cooperation among the majority of nodes is successful. However, the size of the emerged clusters are very different: the firing scenario generates a giant component that is twice as large as the other scenario. The rest of the community apart from the giant component, is located other much smaller clusters. Their size is in the range of [2-12] nodes.

In Figure 3 we investigate how the network topology is close to a *small-world* one. We adopted the small world quotient in order to answer this question. The thick horizontal line represents the threshold value ($> 5$) of $Q$ over which a network is considered small world. Along with $Q$, we also plotted (over the x-y2 axes) the average node productivity to highlight any correlation among them. $Q$ is computed as follows: $Q = \frac{CC_r}{PL_r}$, where $CC_r$ is the ratio among the actual graph clustering

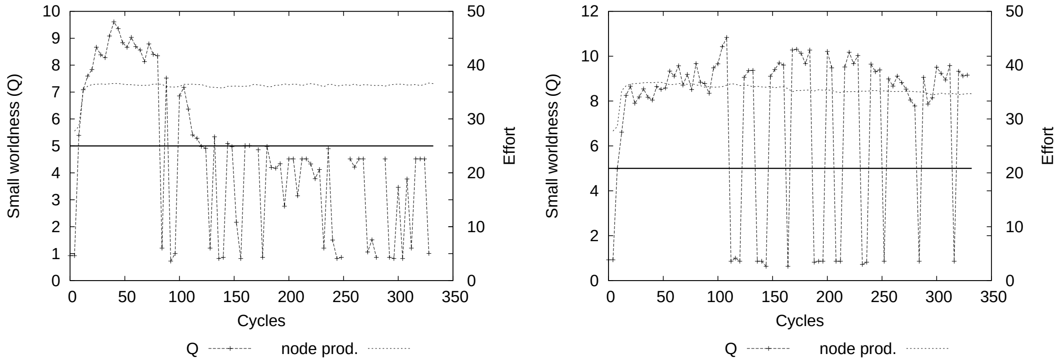


Figure 3: Small worldless $Q$ of , respectively, nofiring and firing scenarios. The thick horizontal line represents the threshold value of $Q$ over which a network is considered small world. Network size is 500.

coefficient and a random graph one ($CC_r = \frac{CC}{CC_{rnd}}$)[1], while $PL_r$ is the ration among the actual graph average path length and a random graph one ($PL_r = \frac{PL}{PL_{rnd}}$)[2]. Our system looks very unstable from this point of view. The nofiring scenario, after cycle 130, has $Q$ always less than the threshold value 5; therefore the network cannot be considered a small world structure. In fact, the visualization of the network - not shown here - depicts clearly the presence of low density structures, similar to sparse trees. The firing scenario is even more unstable. After cycle 100, the rewiring process makes dramatic mutations to the network and it oscillates between a strong small world characterization and a non existent small world characterization; essentially, the network is subjected to random oscillations. These oscillations are easy to be achieved. As the network is composed of sparse tree structures, when new nodes are injected and randomly rewired, there are chances that these newcomers are linked with many distinct trees. This mechanism promotes the newcomers to the role of central hubs, affecting very much the small world coefficient.

[1] We estimated $CC_{rnd}$ using the expression: $N^{-1}$, where $N$ is the network size.

[2] We estimated $PL_{rnd}$ using the expression: $\frac{\log N}{\log k}$, where $N$ is the network size and $k$ is the number of links.

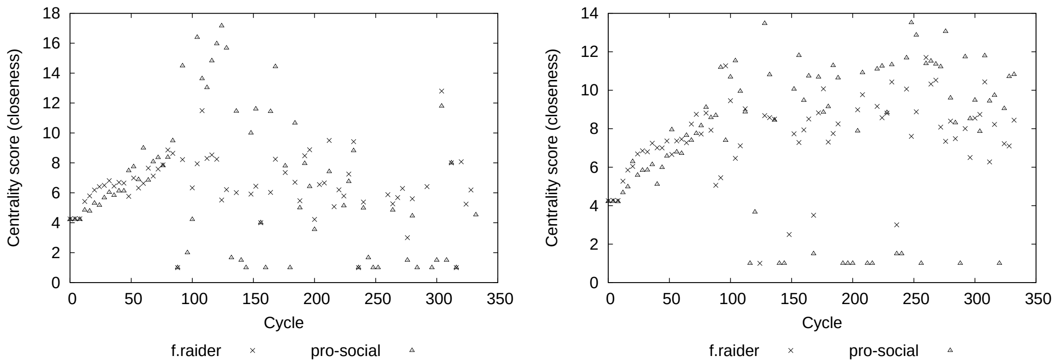


Figure 4: Closeness score among the pro-social and free raider nodes; the left picture shows the nofiring scenario, while the right picture shows the firing one. Network size is 500.

Figure 4 shows the closeness centrality score about the topological location of each node. In general, centrality measures give an indication of the "power“ a node has in the network. We measured the closeness in terms of hop distances. The plots distinguish among pro-social and free raider node scores. Essentially, the closeness centrality aims to emphasizes the distance of a node all other in the network by calculating the distance (in *geodesic path* terms) from each node to all the others. In both cases, we cannot see a significant difference between the closeness properties of pro-social and free raider nodes. This means that they tend to achieve an equal "power" in terms of topological location. This situation is also emphasized by our network model as the CoopNet degree cannot be greater than five. Thus, it is hard to see the emergence of a high connectivity or high density hubs.

Figure 5(a) depicts the system average performance, in terms of effort spent, and Figure 5(b) compares the performance of the firing scenario using distinct firing threshold ($\delta$). These values express the minimum effort tolerated by the system (i.e., $\delta$ = 15, 20 for both fuzzy and strict policy, see Section 6.4). The pictures on the left show the average node's performance, while the pictures on

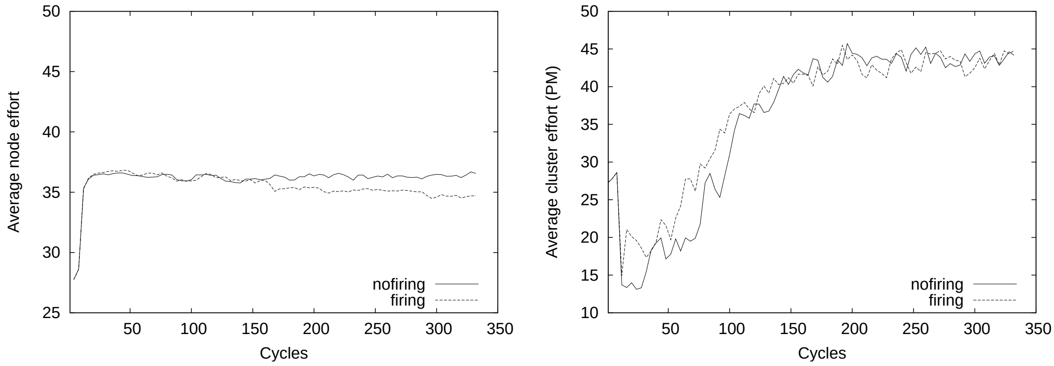


(a) Avg. system effort at node and cluster ($PM$) level

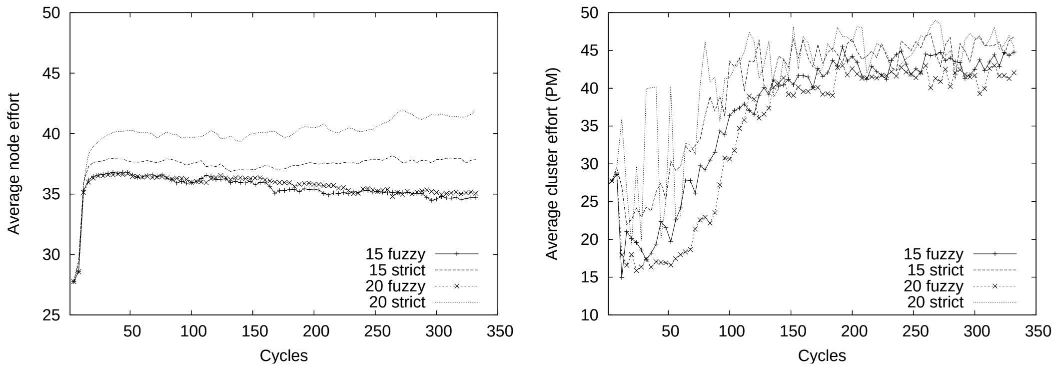


(b) Avg. system effort in firing scenario: distinct firing threshold $\delta$ are compared (with $\delta = 15$, 20 and with strict or fuzzy firing policy)

Figure 5: Comparison among the system average effort. The effort is calculated both at node and cluster level according to the presence or absence of the firing component. These results have been obtained using the $ppm - ww$ rewiring policy. Network size is 500.

the right show the average cluster performance (i.e., $PM_i$).

In general, the system shows an improvement in the average effort pro-fused by the participants since the start of the simulation.

We first focus our attention on Figure 5(a), where the firing threshold $\delta$ is set to 15 using the fuzzy policy. The node performance increases exponentially in the first 25 cycles and than remains almost constant; however, the non-firing scenario shows a better performance than the firing one as it seems to degrade its performance over time. This is an interesting property as it seems that the firing process introduces a sort of *network shock* or *noise* that cannot be tolerated by the system after the middle of the simulation.

While the average performance at the node level highlights the limit of the firing mechanism, at the cluster level (right picture) the performance achieved by the two approaches is almost identical. This is motivated by the different effort distribution that emerges among the cluster participants. Essentially, there are many small groups with very high effort values that bring up the average.

In Figure 5(b) we investigate the impact on performance of distinct firing thresholds. The fuzzy thresholds achieve an almost coincident degrading performance. In both cases the number of nodes affectively removed is very low. Using the strict policy values instead, the system achieves a significant performance improvement at the node level which is proportional to the firing threshold adopted. The improvement is reflected at the cluster level as well, but it is not so evident.

Again, the fuzzy firing mechanisms are the lower performing especially at the beginning of the simulation, while the $\delta = 20$ strict mechanism is still the best, but it produces enormous oscillations. After about 150 cycles, the cluster performance is quite uniform regardless the firing policy adopted. The giant cluster is bigger than in the other scenario (i.e., non firing) and thus less selective because the firing process *lowers the system's entropy* by reducing the diversities among the efforts; essentially, the effort differential is the engine of our system: a lower differential corresponds to a weaker chance to achieve a high effort level. The effort differential becomes smaller and smaller over time and this leads to a uniform performance. Therefore, the cluster performance is a tradeoff among centralized selection (top/down - driven by the firing process) and local selection (driven by individual nodes).

Figure 6 shows the biggest cluster (effort) performance as a function of the concentration of pro-social nodes; the measurement has been taken at the end of

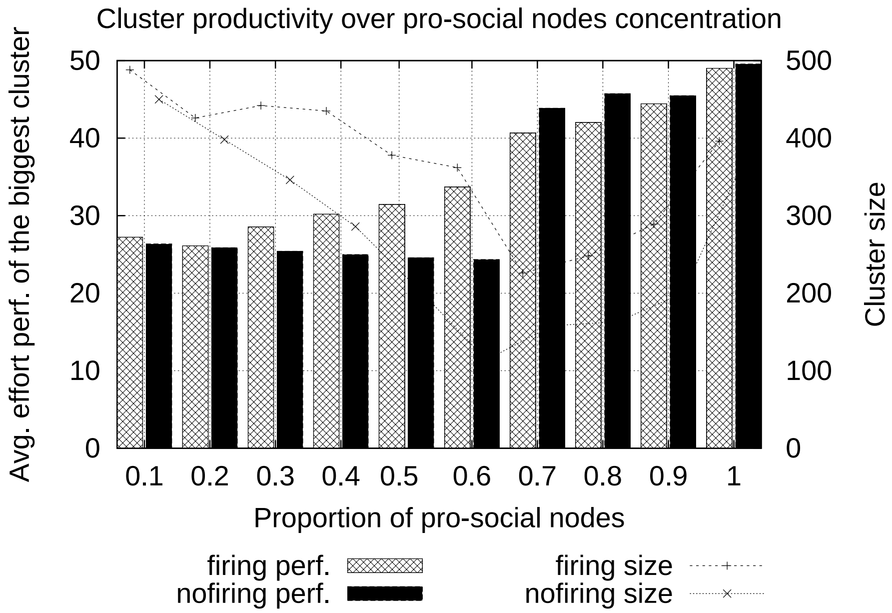


Figure 6: Comparison of the biggest cluster (effort) performance as a function of the concentration of pro-social nodes. On the right axis is represented the biggest cluster size. Both firing and nofiring version are shown. Network size is 500.

each simulation. The results for the firing and nofiring version of the system are compared. In addition, the two dotted lines represent the actual biggest cluster size and their value is reported on the right y axis.

The performance of the firing approach has an almost linear improvement as the concentration of pro-social nodes increases. Surprisingly, using the nofiring mechanism the performance drops slightly, until the concentration of pro-social nodes reaches 0.7; starting from this point, the nofiring seems to be the best.

It is interesting to note that when the proportion of pro-social nodes is very low, the biggest cluster performance is not too bad and, of course, its size is quite close to the entire network. In fact, the free raiders never drop their neighbors and can just lower their effort; therefore the result is an almost non-dynamic network in which the node's effort is averaged towards the average initialization value in a decentralized fashion.

Only when the concentration of pro-socials becomes significant (e.g., from 0.6, 0.7 proportion for, respectively, nofiring and firing mechanisms), the cluster size stops falling, because the pro-socials can satisfy their selective behavior and they do not leave the biggest cluster, boosting its performance.

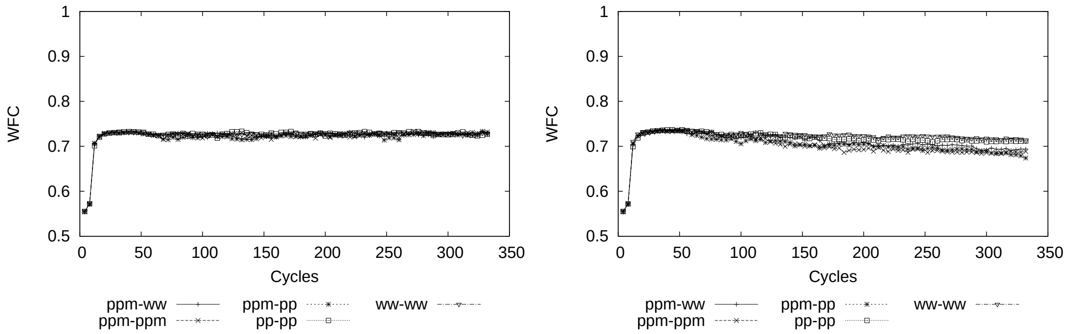


(a) WFC performance in, respectively, nofiring and firing scenarios

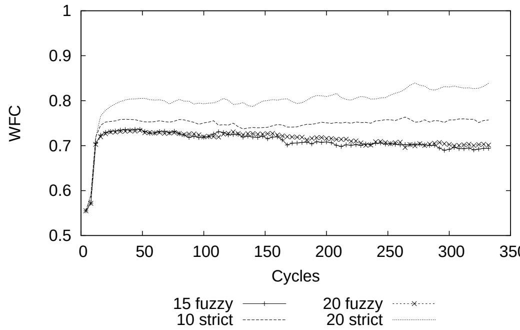


(b) WFC performance in $ppm - ww$ rewiring policy and $\delta = 15, 20$ with both fuzzy and strict firing policies

Figure 7: Welfare From Cooperation(WFC) of, respectively, nofiring and firing scenarios (a). In addition, the WFC performance using distinct firing thresholds $\delta$ is shown in (b). Network size is 500.

We checked the performance of our system using two other metrics: *Welfare From Cooperation* (WFC) and *Marginal Contribution to Cooperation welfare* (MCC). The former gives an indication of the level of welfare originated by the cooperation among nodes and the latter expresses the marginal contribution that each node produces in the system according to its starting effort. These metrics are respectively calculated as follows:

$WFC_i = \frac{p_i}{SYS\_MAX\_EFFORT}$, $WFC = \frac{\sum_{i=1}^{N} WFC_i}{N}$

$MCC_i = \frac{p_i}{p_i(t_0)}$, $MCC = \frac{\sum_{i=1}^{N} MCC_i}{N}$

Figure 7(a) shows the system WFC performance in, respectively, nofiring and firing scenarios. Each line represents a distinct rewiring strategy (see Table 2). All these strategies show a very similar performance. In the first scenario, the WFC level achieves 0.73 and remains constant, showing that the cooperation emerging from the systems works quite well. However, In the second scenario the performance reaches the same level in the same amount of time (about 25 cycles), but then it starts to decrease, regardless the actual rewiring strategy adopted. The top level of $WFC$ is achieved by the $ww - ww$ strategy, which is the least selective policy.

This result is surprising for us as it suggests that firing nodes (i.e., employees) is less efficient that let them self-organize. However, Figure 7(b) tells us that this is not completely true and it could be caused by a bad choice of the firing threshold. In fact, the figure shows the firing scenario $WFC$ performance according to distinct firing threshold (i.e., $\delta = 15$, 20 for both fuzzy and strict policies). Similarly to the average node effort (see Figure 5), using $\delta = 20$ and the strict policy, it achieves the best result. This result is expected as the average node effort and the $WFC$ are tightly correlated.

It is hard to understand deeply why the fuzzy firing strategy performs badly and less than the nofiring one, but at least informally, the idea behind it is that it is better not to remove any node rather than removing a few poor performing ones. The implication of this simple idea is that if we really want to fire nodes and to achieve good performances, we must act with high accuracy. In a real world organization this is not a trivial task and it becomes more and more complicated when the organization structure becomes fluid.

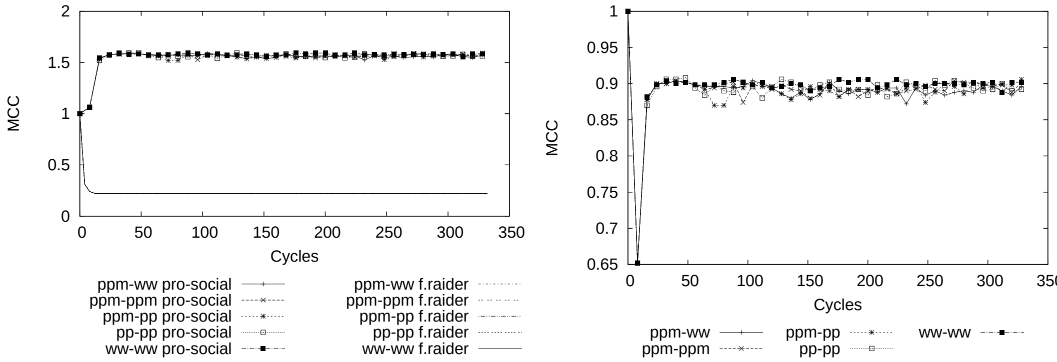


(a) $MCC$ performance, nofiring scenario

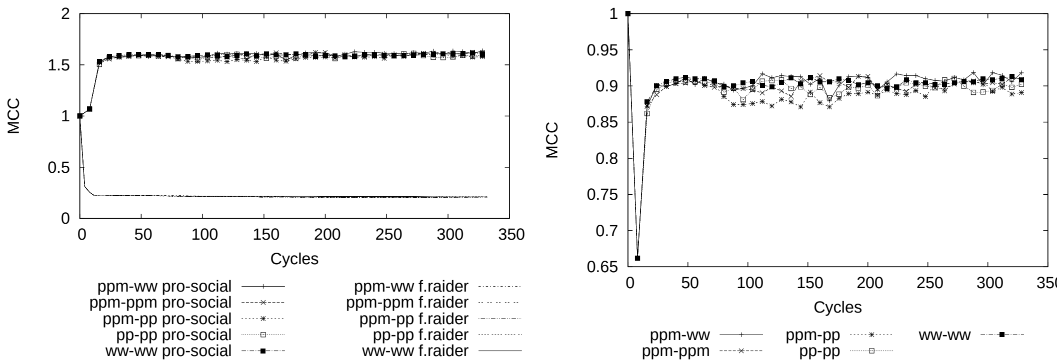


(b) $MCC$ performance, firing scenario

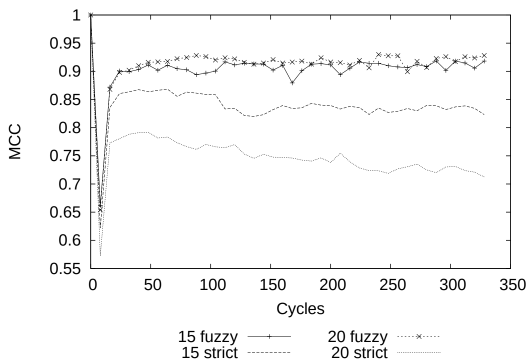


(c) $MCC$ performance in $ppm - ww$ rewire policy according to distinct firing thresholds $\delta = 15$, 20 with both fuzzy and strict firing policies

Figure 8: Marginal Contribution to Cooperation welfare ($MCC$) of, respectively, nofiring (a) and firing (b) scenarios; the results for five distinct rewiring strategies are shown (see Table 2). The pictures on the left distinguishes the $MCC$ performance among pro-social and free raider nodes, while the pictures on the right show the overall $MCC$ performance. Finally, the $MCC$ performance using distinct firing thresholds $\delta$ and policies (fuzzy and strict) is shown in (c). Network size is 500.

The $MCC$ performance is shown is Figure 8. Both nofiring (Figure 8(a)) and firing (Figure 8(b)) scenarios are shown. The plots on the left distinguish among pro-socials and free raiders for each rewiring strategy (see Table 2), while the pictures on the right show the node's overall $MCC$. The $MCC$ level is very good for pro-social nodes no matter the rewiring strategy adopted. However, we expected this result from pro-socials as they can never decrease their initial effort. On the other hand, free raiders achieve a poor performance on average: $MCC = 0.23$. A stable $MCC$ value is achieved quickly, in about 25 cycles, in both scenarios. Selfish nodes can only reduce their effort, but they have the chance to recover their effort to the $p_i(t_0)$ level (initial effort). However, this behavior has no impact on average.

The firing scenario (Figure 8(b)) depicts almost the same result. The only exception is that from cycle 120 the rewiring strategy *ppm-ww* shows a better performance.

In Figure 8(c) we check the effects of the firing threshold on the $MCC$ performance. In contrast with the $WFC$, the $MCC$ depicts the opposite behavior: the lowest threshold, the highest performance. This result is a side effect of the entropy reduction performed by the firing process and of the definition of the $MCC$. As the effort differential becomes narrower, the node's efforts will not move very much from their initial value, therefore the $MCC$ becomes lower because it is defined as the ratio $\frac{p_i}{p_i(t_0)}$.

Figure 9 focuses on the nature of the giant component in both nofiring and firing scenarios. The statistics depict how the effort is distributed according to the degree in the biggest connected part of the network for each available behavior. As expected, pro-socials have a higher $effort/degree$ ratio and they concentrate their effort on a small number of links (2.5 on average). Surprisingly, also free raiders have the same low average degree on which they concentrate their smaller

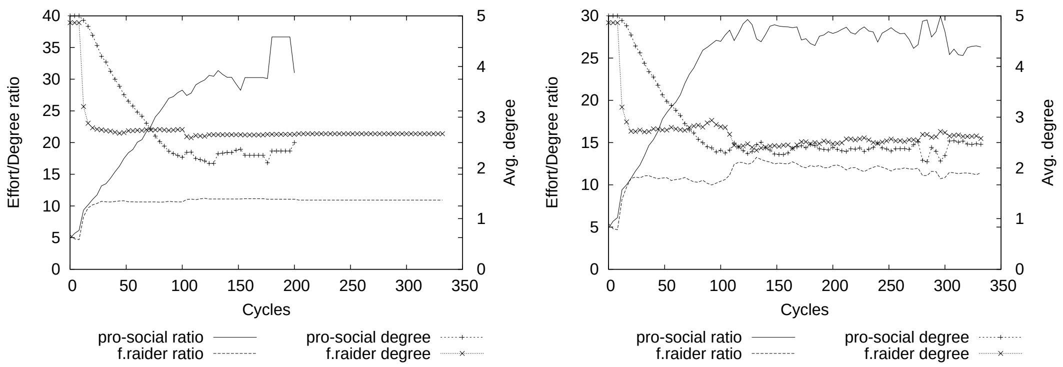


Figure 9: Statistics regarding the giant component. The left picture shows the nonfiring scenario, while the right picture shows the firing one. Network size is 500.

effort.

In the nofiring scenario, all pro-socials leave the giant component, this is why their plots on the graph are interrupted at cycle 200. In addition, when their $effort/degree$ increases[3], their degree remains constant, showing that they are spending more effort; on the contrary, when the $effort/degree$ ratio of pro-socials decreases, their degree increases. Therefore, their effort remains constant probably because on average pro-social nodes perform much better and they may have reached the $SYS\_MAX\_EFFORT$ value.

The E/I-Index is a measure of the connection of a network. It is computed, for each imposed partition (group), as the ratio among the number of links with external nodes and the number of links with internal nodes, divided by the total number of links. It works in an undirected manner. The partitions we have imposed are achieved according to specific attributes. In our particular case, we adopted as attributes, respectively, the *effort* ($E/I_{eff}$) and the node *behavior* ($E/I_{beh}$). The idea about this measure is to check if the attribute is correlated with the existence of the imposed groups in the network.

[3]We remind that a free raider $i$ can increase its effort to accommodate a newcomer, when its effort is below $p_i(t_0)$.

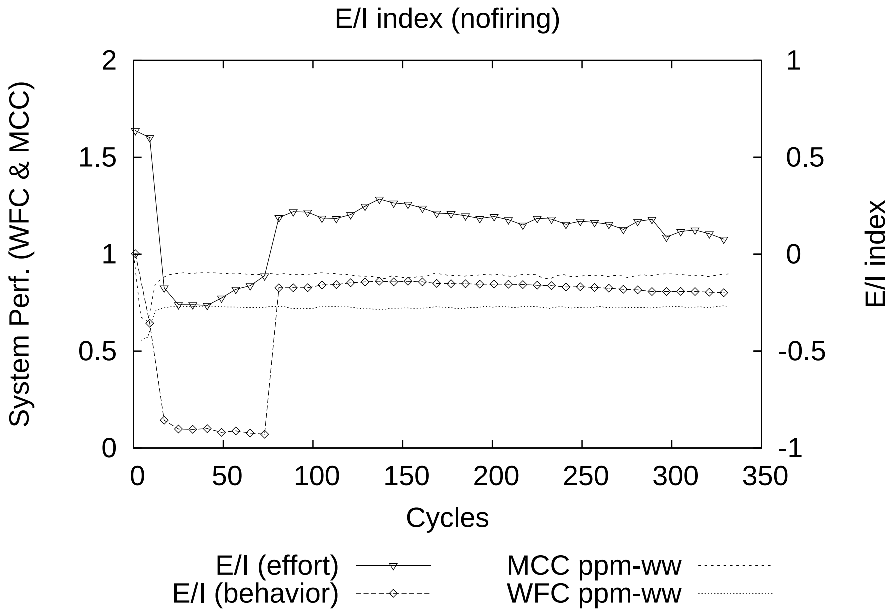


Figure 10: The E/I-Index, the $MCC$ and $WFC$ ratios are compared to highlight any relation among the network group structure and the system performance. Two distinct E/I-Index are considered, respectively effort and behavior based. Network size is 500.

In addition, we checked if the I/E-Index is related to the $MCC$ and $WFC$ performance of the system over time (using *ppm-ww* rewiring strategy, see Section 6.3). The results are shown in Figure 10. The left $y$ axis shows the $MCC$ value, while the right $y$ axis shows the E/I-Index.

At the beginning of the simulation, there is a clear correlation among the E/I, $MCC$ and $WFC$. The $E/I_{eff}$ starts at about 0.6 but drops quickly to a sub zero level (-0.38), indicating that groups based on effort tend to have external links.

The correlation with $WFC$ and $MCC$ is even more clear for the $E/I_{beh}$. Suddenly, at cycle 77 the behavior based groups disappear and the $E/I_{eff}$ assumes a value of about 0.14, showing no correlation among the attribute and the group existence.

At the same time, the trend of $E/I_{eff}$ changed as well. Its value became positive showing that nodes have a light tendency to keep links among the same effort group. However, the index value is unstable and tends to drop and to reach a value too close to zero; this may indicate that the correlation is not very strong.

Essentially, apart from the very beginning, the network organization is not related to groups based on effort classes and on behavior. The network produces its performance (in terms of $WFC$ and $MCC$) in a few cycles (about 25) and in this phase we clearly recognize the partition of pro-social and free raider nodes, then it looks like the network evolves in a random way and the, but this does not affect the performance because the entropy has been already relaxed.

## 7.1 Variable behavior

In the previous sections the proportion among pro-social and free raider nodes was fixed and defined at the beginning of the simulation. Here instead, we still start the experiment with the same 50-50% behavior proportion, but we let each node to change behavior according to a specific policy. The idea is that each node tries to mimic the behavior of its *best* current neighbor, hoping to improve its performance. The notion of *best neighbor* may vary according to the actual policy adopted. We designed two distinct policies in order to select the best neighbor: given a node $i$ and a neighbor $j$, node $j$ is the best neighbor if:

- **w policy**: $j$ has the highest $w$ (wage class) value among the neighbors of $i$ and $w_i < w_j$
- **wrnd policy**: $j$ has the highest $w$ (wage class) value among the random neighbors of $i$ and $w_i < w_j$; the random neighbors are picked from the underlying network graph $G_{rnd}$

As soon as a node detects its best neighbor, it adopts (copies) the neighbor's behavior. The adoption of the possibly new behavior is a stochastic process with probability $p_{bcopy} = 90\ \%$ . In the design of the best neighbor policies we cannot use the $PM$ because two nodes having distinct $PM$ values lie in distinct clusters and cannot be neighbors by definition (in this overlay level).

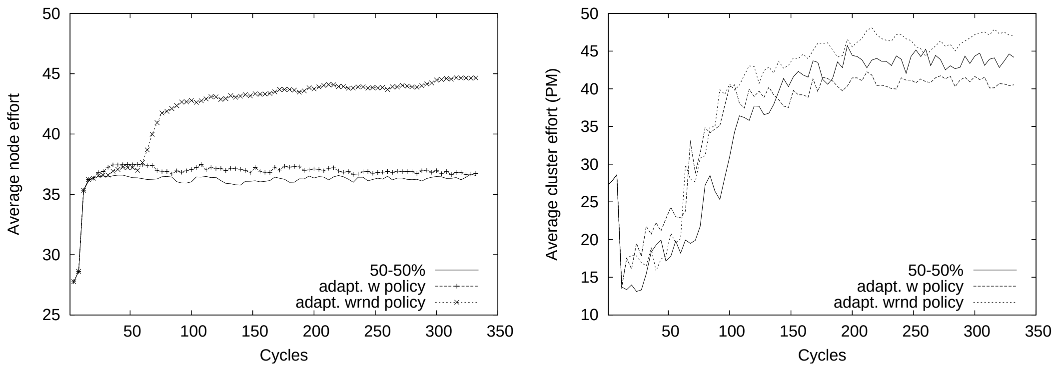


Figure 11: Average node effort and cluster effort performance in the variable behavior scenario. Network size is 500.

Figure 11 shows the average node effort and cluster effort for the variable behavior approach. We only show the result for the nofiring scenario as the firing ones are very similar. The solid line shows the standard 50-50% static proportion of the node's behavior, while the line-points plots depict the w and wrnd policies. The w policy achieves a better performance at the node level, but not at the cluster one. The wrnd policy instead shows a great performance at the node level and it is still the best at the cluster level. This result can be explained in two ways: (i) the degree of the $G_{rnd}$ graph is four times larger than the CoopNet degree, (ii) this set of neighbors is random, allowing each node to contact any peer in the network regardless of its effort or any other CoopNet related parameter. This feature helps enormously the seeking of a suitable neighbor.

Figure 12 shows the $WFC$ performance in the variable behavior scenario. The pictures on the left depict the situation using the nofiring approach, while the pictures on the right depict the firing one. As the $WFC$ is tightly related to the node's effort, we expect a similar performance and in fact the figures are

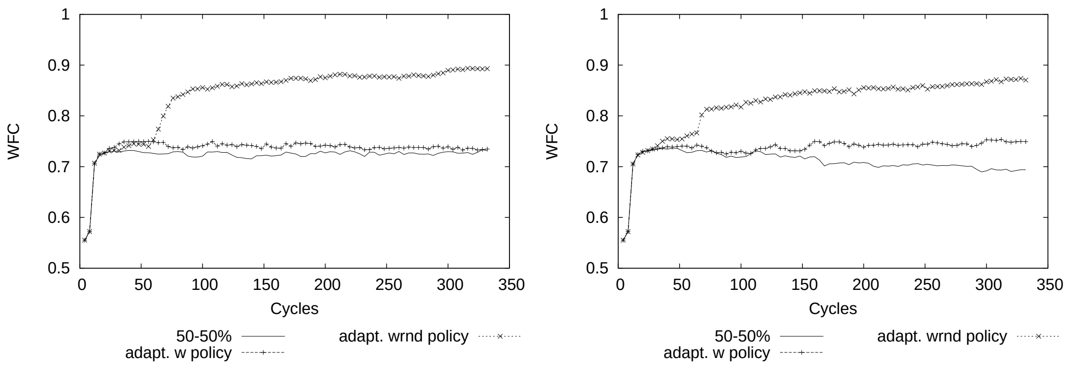


Figure 12: Variable behavior $WFC$ performance summary; the left picture shows the nofiring scenario while the right picture depicts the firing one. network size is 500.

similar. The performance achieved by the variable behavior approach performs better than the previous static 50-50% behavior proportion. In particular, the wrnd policy achieves a very high value, close to 0.9. The performance is almost the same in both firing and nofiring scenarios and it does not tend to degrade as in the static behavior approach regardless the inefficient firing threshold adopted in the experiment ($\delta = 15$ and using the fuzzy policy).

We decided to not include the results for the $MCC$ in this particular scenario. The reason is that the calculation of the $MCC$ value is biased by the $p_i(t_0)$ that changes over time when node $i$ mutates behavior. Therefore, these results are not comparable with the previous ones.

Figure 13 shows the proportion of pro-social nodes that emerges over time in the variable behavior approach. Basically, using the w policy, the proportion of nodes does not change much but oscillates in the [0.53, 0.46] range; this means that periodically each node switch behavior and the system is not stable from this point of view. The wrnd policy instead oscillates much less and leads to the spontaneous emergence of quite a large proportion of pro-social nodes (i.e., 0.63), especially when the firing mechanism is present.

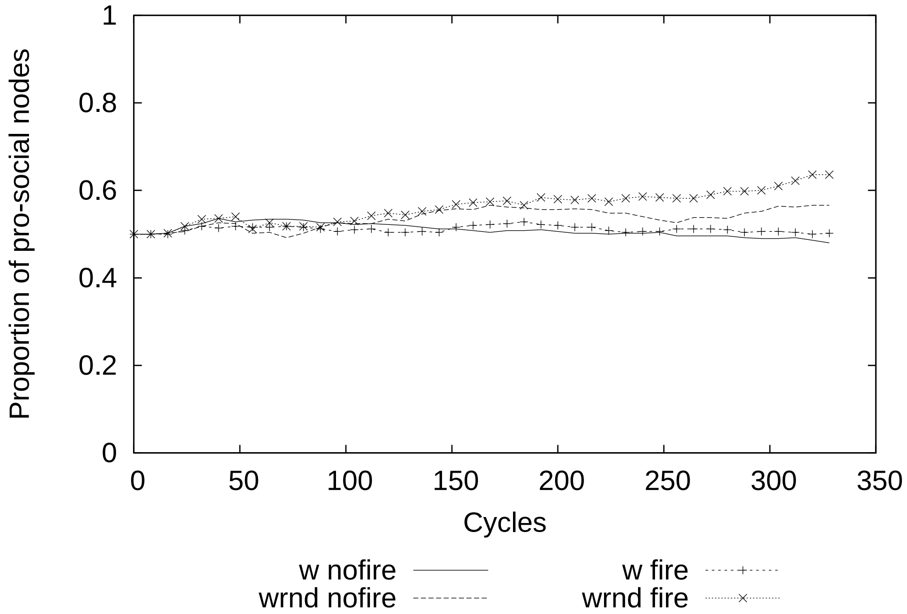


Figure 13: Proportion of emerging pro-social nodes over time. Network size is 500.

Figure 14 compares the WFC performance using distinct network sizes ranging from 500 - our usual size - to 8000 nodes. Both w and wrnd policies are considered. The rewiring policy adopted is $ppm - ww$. Two groups of plots lines emerge from the picture and they correspond to the distinct policies. The network size has no impact on the performance and on time required to reach the corresponding performance level; therefore, the CoopNet approach is *scalable* as the network size and the time required to reach to highest performance are basically not related.

# 8 Conclusions and future works

A number of studies have addressed the problems of social dilemmas and the circumstances under which cooperation among economic agents emerges. Most of these studies have focused upon cooperation and free-riding in generic socio-economic systems; the issue of emerging cooperation within firms, which are specific socio-economic systems, remains an under-research topic.

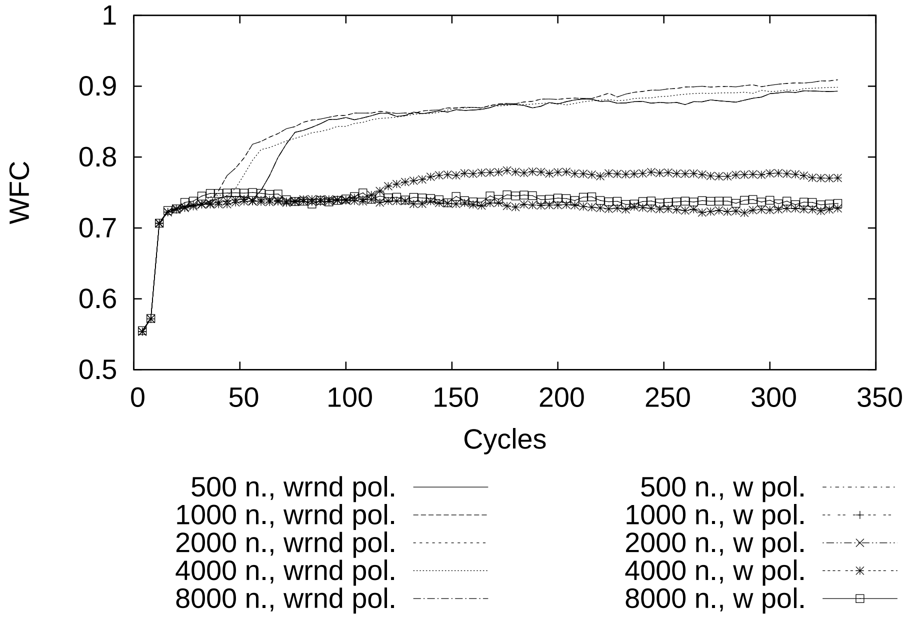


Figure 14: Comparison among the WFC performance using distinct network sizes, ranging from 500 to 8000 nodes. Both w and wrnd policies are shown. The rewiring policy is $ppm - ww$.

The issue of cooperation within firms entails a number of interesting deviations from traditional body of studies on emerging cooperation.

First, within firms, a central hierarchical authority may fire agents who behave as free riders. Second, while studies on emerging cooperation, generally, impinge upon the assumption that individuals are able to confront performances of different strategies and imitate strategies producing higher performances, this process is not straightforward within organisations. Assuming a hierarchical authority that centrally assigns salaries and rewards, not necessarily this authority is able to assign reward as a function of individual effort thereby signalling that it is worth imitating an individual committing high effort to the firm. It is more likely that the central authority is able, at best, to calibrate reward as a function of recognisable teams of individuals within which, however, individuals committing different efforts and adopting different strategies may coexist.

Our simulation experiments highlight two areas of problems in using hierarchical control within an intra-organisational network.

First problem concerns the relationship between firing and self-organisation of intra-organisational networks. Simulations show that hierarchical intervention in removing nodes that are social loafers is appropriate when it is possible to select precisely which nodes are underperforming. It is better not to remove any node rather than removing a few very poor performing ones. Thus is true because firing brings about a high social cost in term of network evolution.

On one hand, firing interrupts longitudinal exchange behaviour, which is mechanism that allows pro-social nodes to learn and scale up their effort; on the other hand, by eliminating the low performing nodes, it enforces a more uniform effort distribution (i.e., less entropy) among the participants. This distribution allows the participants to find more easily an equilibrium state and many more links (relations) are successfully maintained. This is the reason why, with firing, the larger component is twice as in the simulations without firing. However, this does not automatically translates into better performance. On the contrary, when the firing is stochastic, there are good chances that very poor performing nodes are not purged from the system. This fact combined with the lowered effort entropy has negative effect on the CoopNet performance. In fact, it would be much better to let the system self-organize in this case.

The implication of this simple idea is that if we really want to fire nodes and to achieve good performances, we must act with high accuracy. In a real world organization this is not a trivial task and it becomes more and more complicated when the organization structure becomes fluid.

Second problem concerns the interaction between reward structure and imitation policies. In general, our simulations, show that imitation of best performing agents improve performances. However, this is true only if imitation regards nodes that are not connected to the neighbourhood that embeds the individuals that is imitating. Indeed, if reward is assigned to recognisable groups, the closer the individual imitated, the higher is the possibility that the salary does not incorporate valuable information. In our simulations, imitation policies are effective when nodes imitate other individuals that are randomly selected. When they imitate other individuals either in their neighbourhood or connected to their neighbourhood, it is unlikely that the imitated individual belongs to another group, which is recognisable by a central authority as eligible for a different reward. As a consequence, such an imitation behaviour does not convey interesting information to direct strategy selection.

# References


[1] Peersim: a p2p simulator.

[2] P. S. Adler. Market, hierarchy, and trust: The knowledge economy and the future of capitalism. *Organization Science*, 12:215–234, 2001.

[3] R. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.

[4] R. Axelrod. Advancing the art of simulation in the social sciences. *Complex.*, 3(2):16–22, 1997.

[5] K. R. Conner and C. K. Prahalad. A resource-based theory of the firm: Knowledge versus opportunism. *Organization Science*, 7:477–501, 1996.

[6] P. Costa, V. Gramoli, M. Jelasity, G. P. Jesi, E. Le Merrer, A. Montresor, and L. Querzoni. Exploring the Interdisciplinary Connections of Gossip-based Systems. *Operating System Review*, 41(5):51–60, October 2007.

[7] V. Eguiluz, M. G. Zimmermann, C. J. Cela-Conde, and M. San Miguel. Cooperation and the emergence of role differentiation in the dynamics of social networks. *American Journal of Sociology*, 110(4):977–1008, 2005.

[8] J. Epstein and R. Axtell. *Growing Artificial Societies: Social Science From The Bottom Up*. Brookings Institute Press, 1996.

[9] N. J. Foss. *Strategy, Economic Organization, and the Knowledge Economy.* Oxford University Press, 2005.

[10] A. Gouldner. The norm of reciprocity: a preliminary statement. *American Sociology Review*, 47:73–80, 1960.

[11] R. M. Grant. Toward a knowledge-based theory of the firm. *Strategic Management Journal*, 17:109–122, 1996.

[12] D. Hales, B. Edmonds, E. Norling, and J. Rouchier, editors. *Multi-Agent-Based Simulation III, 4th International Workshop, MABS 2003, Melbourne, Australia, July 14th, 2003, Revised Papers*, volume 2927 of *Lecture Notes in Computer Science*. Springer, 2003.

[13] W. D. Hamilton. The genetic evolution of social behavior. *Journal of Theoretical Biology*, 7:1–52, 1964.

[14] N. Hanaki, A. Peterhansl, P. S. Dodds, and D. J. Watts. Cooperating in evolving social networks. *Management Science*, 53(7):1243–1248, 2007.

[15] G. Hardin. The tragedy of the commons. *Science*, xx:1243–47, 1968.

[16] G. M. Hogdson. *Economics and Utopia*. Routledge, 1999.

[17] M. Jelasity, R. Guerraoui, A.-M. Kermarrec, and M. van Steen. The peer sampling service: Experimental evaluation of unstructured gossip-based implementations. In *Middleware*, Oct. 2004.

[18] M. Jelasity, W. Kowalczyk, and M. van Steen. Newscast computing. Technical Report IR-CS-006, Vrije Universiteit, Amsterdam, The Netherlands, Nov 2003.

[19] D. Lepak and S. Snell. The human resource architecture: Toward a theory of human capital allocation and development. *Academy of Management Review*, 24(1):31–48, 1999.

[20] B. Lomborg. Nucleus and shield: The evolution of social structure in the iterated prisoner's dilemma. *American Sociological Review*, 61:278–307, 1996.

[21] O. M. *The Logic of Collective Action*. Harvard University Press, Cambridge, MA, 1965.

[22] G. MacCormack. Reciprocity. *Man*, 11(1):89–103, 1976.

[23] M. Macy. Learning to cooperate: Stochastic and tacitt collusion in social exchange. *American Journal of Sociology*, 97(3):808–843, 1991.

[24] M. Macy and J. Skvoretz. The evolution of trust and cooperation between strangers: A computational model. *American Sociological Review*, 63(5):638–660, 1998.

[25] M. Mauss. *The Gift: The Form and Reason for Exchange in Archaic Societies.* W. W. Norton & Company, August 2000.

[26] R. M. May, S. Bonhoeffer, and M. A. Nowak. Spatial games and evolution of cooperation. In *Proceedings of the Third European Conference on Advances in Artificial Life*, pages 749–759, London, UK, 1995. Springer-Verlag.

[27] M. Nettle and R. I. M. Dunbar. Social markers and the evolution of reciprocal exchange. *Current Anthropology*, 38(1):93–99, 1997.

[28] M. A. Nowak and R. M. May. Evolutionary games and spatial chaos. *Nature*, 359(6398):826–829, October 1992.

[29] A. Rapoport and A. M. Chammah. *Prisoner's Dilemma: A study in Conflict and Cooperation.* Ann Arbor: University of Michigan Press, 1965.

[30] R. L. Trivers. The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46:35–57, 1971.

[31] S. Voulgaris, D. Gavidia, and M. van Steen. CYCLON: Inexpensive Membership Management for Unstructured P2P Overlays. *J. Network Syst. Manage.*, 13(2), 2005.