

Document downloaded from:

<http://hdl.handle.net/10251/40211>

This paper must be cited as:

Del Val Noguera, E.; Rebollo Pedruelo, M.; Botti V. (2014). Strategies for cooperation emergence in distributed service discovery. *Cybernetics and Systems*. 45(3):220-240. doi:10.1080/01969722.2014.894848.



The final publication is available at

<http://dx.doi.org/10.1080/01969722.2014.894848>

Copyright Taylor & Francis: STM, Behavioural Science and Public Health Titles

Strategies for cooperation emergence in distributed service discovery

E. del Val, M. Rebollo, V. Botti

Departament de Sistemes Informàtics i Computació,
Universitat Politècnica de València,
Camí de Vera s/n. 46022, València. Spain
{edelval,mrebollo,vbotti}@dsic.upv.es

Abstract. In distributed environments where entities only have a partial view of the system, cooperation plays a key issue. In the case of decentralized service discovery in open agent societies, agents only know about the services they provide and who are their direct neighbors. Therefore, they need the cooperation of their neighbors in order to locate the required services. However, cooperation is not always present in open systems. Non-cooperative agents pursuing their own goals could refuse to forward queries from other agents to avoid the cost of this action; therefore, the efficiency of the decentralized service discovery could be seriously damaged. In this paper, we propose the combination of incentives and local structural changes in order to promote cooperation in the service discovery process. The results show that, even in scenarios where the predominant behavior is not collaborative cooperation emerges.

1 Introduction

There are distributed systems where the cooperation of all the entities that participate in them is required to obtain a good performance that provides benefits for all the participants. Some of the scenarios where cooperation is required are: wireless ad-hoc networks where nodes rely on other nodes to forward their packets in order to reach the destination node; file sharing in P2P systems [1]; streaming applications [2], discussion boards [3], on-line auctions [4], or overlay routing [5].

If participants do not contribute in order to maximize their own benefits and exploit the contributions of the others, they will obtain a high rate of benefits in the short term. However, these benefits decrease as the number of selfish participants increases, thereby damaging the performance of the whole system. There are models of genetic and cultural evolution that confirm that the opportunity to take advantage of others undermines and often eliminates cooperation [6]. These cooperation problems are also known as social dilemmas (i.e., the tragedy of the commons, the free-rider problem, the social trap). The promotion and stabilization of cooperation in these scenarios has been considered to be an area of interest [7].

In this paper, we present a proposal that promotes cooperation in the service discovery process among agents that are located in a network structure. In this context, cooperation plays an important role since agents only have a partial view of the network and need the cooperation of their neighbors to find the required provider agent.

This becomes even more difficult when there are self-interested agents that do not cooperate with other agents in order to avoid the cost of forwarding queries. Our proposal combines incentives and social plasticity to promote cooperation during the service discovery process.

The paper is structured as follows. In section 2, we present other works related to cooperation emergence. In section 3, we describe the model where we integrate the cooperation mechanisms. This section contains the description of the service discovery process, presents the incentives mechanism and the social plasticity, and finally we describe how agents select each action during the service discovery process. Section 4 presents a set of experiments where we evaluate the performance of our proposal. Finally, section 5 presents conclusions and final remarks.

2 Related Work

Several mechanisms have been proposed to promote and maintain cooperation in different scenarios. In scenarios where individuals interact repeatedly, selfish or altruistic actions would be returned in future. Therefore, a common mechanism to facilitate the emergence of cooperation is *direct reciprocity* [8]. When agents do not always interact with the same individuals, *indirect reciprocity* [9] or *tags* [10] are used. *Punishment* has also been considered to promote cooperation and to overcome the "tragedy of the commons" [6]. Punishment is present in human societies where sanctioning institutions apply a punishment to those that do not obey the law. In systems where such centralized institutions do not exist, individuals are willing to punish defectors even though this implies a cost for them [11]. In general, punishment has been proven to be an efficient way to maintain cooperation [12, 13].

Many approaches that are used to promote cooperation assume well-mixed populations where everybody interacts with equal frequency with everybody else. However, real populations are not well-mixed. In real scenarios, some individuals interact more often than others; therefore, to understand the social behavior of the systems it is important to consider the social structure. The social structure is represented by a network where links are established by the individuals following certain preferences. There are several works that analyze the influence of the network structure in the emergence of cooperation. These works study how structural parameters such as clustering or degree distribution affect the emergence and maintenance of cooperation [14–17].

Another issue that it is important to consider is how local changes can influence the collective social behavior. Eguíluz et al. [18] present a model that uses the Prisoner's Dilemma game [19] and *social plasticity* in random undirected networks of agents. Agents update their behavior in discrete time steps using an imitation strategy that considers the payoff of neighbors. Agents use social plasticity (i.e., changes in structural links) to facilitate the replacement of an unprofitable relationship with a new one that is randomly chosen. Griffiths et al. [20] propose a mechanism that considers context awareness and tags of agents to promote cooperation. Moreover, agents can remove part of their connections with agents that are not cooperative and add connections with others that can improve cooperation. There are other approaches that also make use of

rewiring techniques and partial observation to facilitate the emergence of cooperation [21].

The majority of the proposals present in the literature considers incentives and local structural changes in the network separately. In this paper, we integrate both mechanisms and analyze the effect of this integration. Specifically, we present a proposal that promotes cooperation in the service discovery process among agents that are located in a network structure. In this context, cooperation plays an important role since agents only have a partial view of the network and need the cooperation of their neighbors in order to forward queries to locate the required provider agents. This becomes even more difficult when there are self-interested agents that do not cooperate with other agents in order to avoid the cost of forwarding queries. We combine two mechanisms to promote cooperation: incentives and social plasticity. The main differences between our approach and other proposals are: (i) we consider the social structure where agents are located instead of a well-mixed population; (ii) we have considered different criteria for the assignment of incentives for the agents that participate in the search process; (iii) local structural changes are also taken into account in combination with incentives; the structural changes are not random, agents break links with those neighbors that have non-cooperative behavior, and instead of replacing them randomly, agents look for another agent based on their preferences; (iv) taking into account local information about the degree of cooperation of their neighborhood, agents are able to detect when it is more appropriate the use of social plasticity in combination with incentives. The proposed mechanisms have been tested and the results show that even in adverse situations where there is a large number of non-cooperative (non-cooperator) agents our proposal obtains good results and the performance of the system is not seriously affected.

3 Model for Cooperation in Service Discovery

Consider a network of agents $A = \{1, \dots, n\}$ connected by undirected links in a fixed network represented by the adjacency matrix g . A link between two agents i and j , such that i and $j \in A$, is represented by $g_{ij} = g_{ji} = 1$, where $g_{ij} = 0$ means that i and j are not connected. The set of neighbors of agent i is $N_i = \{j | g_{ij} = 1\}$. We assume that $g_{ii} = 0$. The number of neighbors of i is denoted by k_i , which is the cardinality of the set N_i (see Figure 1).

Agents in the system are characterized by the roles they play. The organizational role determines the type of services offered by the agent. A role r_i is defined by a semantic concept defined in an organizational ontology, and a set of service semantic descriptions associated to the role. Each service description s_i is defined by inputs, outputs, preconditions, and effects of the services. An agent has an initial *behavior* that can be cooperative (*coop*) or not cooperative (*ncoop*). Moreover, each agent has an initial budget b that it is equal for all the agents in the system.

A link between two agents i and j ($g_{ij} = 1$) is established considering a probability. This probability is based on the similarity between the roles played by the agents i and j and the services provided by them as well as their degree of connectivity. Therefore, agents have a greater probability of establishing links with agents that have similar attributes than with dissimilar ones. The result of using this criterion to establish links

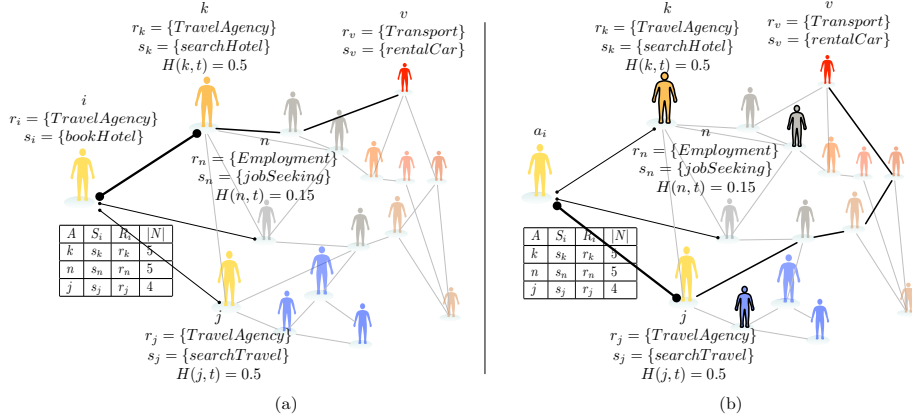


Fig. 1: Service discovery scenario. (a) All the agents are cooperative. Agent i only knows its direct neighbors k , j , and n . If i needs to locate a service (i.e., rentalCar), it will forward the query to its most promising neighbor (i.e., k) based on the homophily between the neighbor and the target agent (i.e., tg) that should provide the required service and the degree of the neighbor. (b) Agents with thick silhouette are not cooperative. Agent i has to find an alternative agent to k since k refuses to forward queries.

between agents is a network structure based on similarity and degree that has an exponential distribution of its degree of connection. This structure facilitates the task of decentralized service discovery only considering local information. For further details about the process of network creation we refer the reader to [22].

3.1 Service Discovery

The service discovery starts when agent $i \in A$ needs to locate an agent that plays certain role and offers certain service in order to deal with one of its goals (see Alg. 1). The agent i , in order to start the process, estimates if it has enough budget b to reach the target. This estimation is based on the information from previous searches. In the case that the budget is enough, agent i creates a query at time t , $q^t = \{i, s_{tg}, r_{tg}, TTL, \varepsilon, \{\}\}$, which consists of: the agent that starts the service discovery process i , the required semantic service description (s_{tg}), the organizational role that the target agent should play (r_{tg}), the Time To Live that represents the maximum number of times that the query can be forwarded (TTL), a similarity threshold ε established by i that represents how similar should be the service offered by an agent to consider that the target agent has been found, and the list of identifiers of the agents that participate in the discovery process (initially this list is empty).

In the discovery process, when an agent that is similar enough to the target is found, the agent i is informed and the process ends (see Alg. 2 Line 6). Otherwise, agent i should choose one of its neighbors to forward the query q^t (see Alg. 2 Line 28). The selection of the the most promising neighbor is based on a probability $P(\langle j, tg \rangle)$ that

Algorithm 1 Function that describes how an agent i starts a service discovery process.

```

1: function startDiscovery()
2:  $TTL \rightarrow 100$  /* number of steps associated to the query */
3:  $path \rightarrow \{\}$  /* list of participants in the service discovery process */
4:  $\varepsilon \rightarrow 0.75$  /* threshold similarity */
5: serviceDiscovery( $i, q = (i, s_{tg}, r_{tg}, TTL, \varepsilon, path)$ )
6: end function

```

considers: semantic similarity and degree of connection. The semantic similarity is calculated between the neighbor and the target (i.e., similarity-based factor that considers the semantic similarity between the services and the roles of two agents) and the degree of connection refers to the degree of connection of the neighbor [22].

For each neighbor j , $P(\langle j, tg \rangle)$ determines the probability that the neighbor j redirects the search to the nearest network community where there are more probabilities of finding the agent tg .

$$P(\langle j, tg \rangle) = 1 - \left(1 - \left(\frac{H(j, tg)}{\sum_{k \in N_i} H(k, tg)} \right) \right)^{k_j} \quad (1)$$

where $H(j, tg)$ is the semantic similarity between the roles and services of agents j and tg , and k_j is the degree of connection of neighbor j . For a detailed mathematical definition of H we refer the reader to [22]. The agent i selects the neighbor $j \in N_i$ that maximizes the probability $P(\langle j, tg \rangle)$.

$$\mathcal{F}_{N_i}(tg) = \operatorname{argmax}_{j \in N_i} P(\langle j, tg \rangle) \quad (2)$$

The discovery process ends when the number of forwards exceeds the TTL or when the target agent that provides the required service is found.

Actions and Incentives. During the service discovery process, when an agent i receives a query q^t , it has to choose an action a_i among a set of possible actions $Acc = \{\rho, \infty, 1, 2, \dots, k_i, \emptyset, \lambda\}$, where:

- ρ is asking for a service
- ∞ is providing the service
- $\{1, \dots, k_i\}$ is forwarding the query to one of its neighbors $\in N_i$
- \emptyset is doing nothing
- λ rewiring a link

These actions have associated a cost, a benefit, or a reward. If an agent asks for a service to a provider, it has to pay the provider β . If an agent provides a service, it earns a payoff p . Forwarding a query is costly c , but an agent earns a payoff α if the query ends successfully. Otherwise, the payoff is 0. If an agent chooses the action \emptyset , its payoff is 0. The agent can also decide rewiring a current structural relation with a neighbor and looking for a new one. The rewiring action has a cost γ . Formally:

Algorithm 2 Function that describes the service discovery process that an agent i carries out when it receives a query from another agent l . The parameters of the function are: l the agent that sends the service discovery request, i the agent that receives the service discovery request and q^t the query

```

1: function serviceDiscovery( $l, i, q^t = (i', s_{tg}, r_{tg}, TTL, \varepsilon, path)$ )
2: /* target agent that represents the profile of a suitable provider agent to solve the query */
3:  $tg \leftarrow (r_{tg}, s_{tg}, \emptyset, \emptyset)$ 
4: if  $TTL > 0 \wedge b > costEstimation()$  then
5:   /* if an agent  $i$  is enough similar to the target */
6:   if  $H(i, tg) \geq \varepsilon$  then
7:     /* agent  $i$  obtains a benefit for offering the service */
8:      $\pi_i^t(behavior) \rightarrow \pi_i^t(behavior) + u_i^t(\infty)$ 
9:     /* asking a service has a cost for the agent  $i'$  */
10:     $\pi_{i'}^t(behavior) \rightarrow \pi_{i'}^t(behavior) + u_{i'}^t(\rho)$ 
11:    for  $n \in path$  do
12:       $\pi_n^t(behavior) \rightarrow \pi_n^t(behavior) + u_n^t(a_n^t = \emptyset \wedge \exists t' \leq t : a_n^{t'} \in \{1, 2, \dots, k_i\} \wedge \exists j \in A : a_j^t = \infty)$ 
13:       $\pi_{i'}^t(behavior) \rightarrow \pi_{i'}^t(behavior) - u_n^t(a_n^t = \emptyset \wedge \exists t' \leq t : a_n^{t'} \in \{1, 2, \dots, k_i\} \wedge \exists j \in A : a_j^t = \infty)$ 
14:      /* sends an inform to the agent  $i'$  that started the process */
15:       $inform(i', i, s_q)$ 
16:    end for
17:  else
18:    /* if the agent is not enough similar */
19:    /* if the agent has a non-cooperative behavior..*/
20:    if  $i$  is ncoop then
21:       $\pi_i^t(behavior) \rightarrow \pi_i^t(behavior) + u_i^t(\emptyset)$ 
22:       $r_{qi} \leftarrow r_{qi} + 1$ 
23:      socialPlasticity( $r_{qi}, \sigma$ )
24:      /* sends a reject message to the agent  $i'$  that started the process */
25:       $reject(l, i, s_q)$ 
26:    else
27:      /* it selects the most promising neighbor for forwarding*/
28:       $j \leftarrow \mathcal{F}_{N_i}(tg)$ 
29:      /* agent  $i$  updates its utility function with the cost of the forwarding action */
30:       $\pi_i^t(behavior) \rightarrow \pi_i^t(behavior) + u_i^t(\{1, 2, \dots, k_i\})$ 
31:       $TTL \leftarrow TTL - 1$ 
32:      serviceDiscovery( $i, j, q^{t+1} = (i', s_{tg}, r_{tg}, TTL, \varepsilon, path \cup \{i\})$ )
33:    end if
34:  end if
35:  /* if the TTL is exhausted */
36:   $inform(i', \emptyset, \emptyset)$ 
37: end if
38: end function

```

$$u_i^t(a_i^t) = \begin{cases} -\beta & \text{if } a_i^t = \rho \\ p & \text{if } a_i^t = \infty \\ -c & \text{if } a_i^t \in \{1, 2, \dots, k_i\} \\ 0 & \text{if } a_i^t = \emptyset \wedge \nexists t' \leq t : a_i^{t'} \in \{1, 2, \dots, k_i\} \\ \alpha & \text{if } a_i^t = \emptyset \wedge \exists t' \leq t : a_i^{t'} \in \{1, 2, \dots, k_i\} \wedge \exists j \in A : a_j^t = \infty \\ -\gamma & \text{if } a_i^t = \lambda \end{cases} \quad (3)$$

where $u_i^t(a_i^t)$ is an utility function that calculates the payoff obtained by an agent i when it executes an action $a_i^t \in Acc$ in time t (see Alg. 2).

Once the service discovery process ends successfully, incentives are distributed among the agents that participated forwarding a query (see Alg. 1). The use of incentives tries to provide a reward the effort to those agents that cooperate during the discovery process. We have considered different types of mechanisms to distribute incentives:

- mechanisms that *uniformly* distribute the incentives among all the agents that participated in the forwarding process of a query that ended successfully. We consider two different mechanisms that differ from who is the entity that provides the incentives to the other agents. In one mechanism the system is the entity responsible of providing the incentives to the agents. This mechanism is called *System*. In the other mechanism, the agent that initiates the discovery process provides the incentives to all the participants. We called this mechanism *Fixed*.
- mechanisms that use a criterion to distribute the incentives in a *non-uniform* way among all the agents that participated in the forwarding process of a query that ended successfully.
 - *Path*: the reward depends on the length of the path. The shorter path to locate the provider agent is, the higher reward the agents will received. This criterion tries to reward agents that are part of short paths.
 - *SimDg*: the reward for an agent that participates in the forwarding process depends on its similarity with the target agent and its degree of connection. The participants that are closer to the target agent (i.e., they are similar to the target) and have a high degree of connection will receive a higher reward than the other participants. This criterion rewards agents that are well connected and close to the target.
 - *InvSimDg*: the reward for an agent that participates on the forwarding process depends on its difference with the target agent and its degree of connection. The participants that are distant to the target agent and have a low degree of connection will receive a higher reward that the other participants. This criterion tries to reward those agents that cooperate although they are distant to the target agent.

Social Plasticity. The structure of the network influences interactions of agents, therefore it is important to provide agents mechanisms to be able of changing their local structure in the network. For that reason, we consider the rewiring action λ in our

model. Through interactions during the service discovery process, agents are able to change their structural relations taking into account which neighbors provide profitable relationships and which do not. This feature is called social plasticity [18]. Social plasticity is the capacity of individuals to change their relationships as time passes. Specifically, in our system, each agent maintains information related to its neighbors. This information consists of the number of times a neighbor $j \in N_i$ has refused to forward one of its queries (rq_{ij}) (see Alg. 2 Line 22-23).

In order to evaluate the utility of a link, an agent i uses a decay function that calculates the probability of maintaining a link with j taking into account the number of queries that it would have sent through neighbor j but j refused to forward. This function is a sigmoid that ranges between $[0,1]$.

$$P_{decay}(rq_{ij}) = \frac{1}{1 + e^{\frac{-(rq_{ij}-d)}{y}}} \quad (4)$$

where the constant y is the slope and d is the displacement. These constants are established by the agent. The most influential constant is d . The displacement d indicates how benevolent an agent is with respect the non-cooperative behavior of its neighbors. A high value of d means that the agent is going to consider a higher number of refuses in order to make a decision about looking for another neighbor. A low value means that it is not permissive with the number of refuses. The function $P_{decay}(rq_{ij})$ returns a value in the range $[0,1]$, where 0 indicates that the agent does not consider that the number of rejects from its neighbor is enough to make a decision about rewiring, and 1 indicates that it is necessary to change the link. If an agent decides to break a link, it looks for a candidate to replace it. The criterion commonly used in other works is random (i.e., an agents selects a random agent to establish a link). However, in our proposal, agents look for a neighbor that offers similar services to the previous neighbor in order to maintain the structure of the network. We assume that agents accept links from other agents since this fact increases their connectivity in the network.

In order to find a trade-off between the number of structural changes and the emergence of cooperation, the use of the rewiring action λ by an agent is affected by the number of cooperator neighbors. If the number of cooperator neighbors is under a certain threshold σ , the mechanism used to facilitate the emergence of cooperation is the social plasticity combined with incentives. Otherwise, the mechanism used is based on incentives only.

3.2 Action Selection

Agents choose which will be the next action taking into account: (i) the similarity between itself and the target agent; (ii) previous actions of their neighbors. An agent i has an information structure $H_i^t = \{\pi_i^t(coop), \pi_i^t(ncoop)\}$ that stores information about the budget that the agent has when its behavior was cooperative $\pi_i^t(coop) = \sum_{t' \leq t} u_i^{t'}(a_i^{t'}), a_i^{t'} \in Acc - \{\emptyset\}$ and when it was non-cooperative $\pi_i^t(ncoop) = \sum_{t' \leq t} u_i^{t'}(a_i^{t'}), a_i^{t'} \in Acc - \{1, \dots, k_i\}$. Moreover, an agent i stores the number of times it sends a query to one of its neighbor j and it rejected forwarding it (rq_{ij}).

When an agent i receives a query q^t at time t , it chooses one of these actions using the following criterion:

- do the task itself when its service and role are enough similar to the service and role of the target agent tg .

$$a_i^t = \infty \text{ if } |H(i, tg)| \geq \varepsilon \quad (5)$$

- do nothing when its service and role are not enough similar to the target agent and, considering information from previous stages, agent i finds that the neighbor with highest benefit did not cooperate in the previous stage $t - 1$.

$$a_i^t = \emptyset \text{ if } |H(i, tg)| < \varepsilon \wedge a_j^{t-1} = \emptyset, j \in \text{argmax}(H_1^{t-1}, \dots, H_{k_i}^{t-1}) \quad (6)$$

- forwarding the query to one of its neighbors $j \in N_i(g)$ when its service and role are not enough similar to the target agent and the neighbor with highest benefit cooperated in the stage $t - 1$.

$$a_i^t = j \text{ if } |H(i, tg)| < \varepsilon \wedge a_j^{t-1} \neq \emptyset, j \in \text{argmax}(H_1^{t-1}, \dots, H_{k_i}^{t-1}) \quad (7)$$

where

$$j \in \text{argmax}_{j \in \{1, \dots, k_i\}} P(\langle j, tg \rangle) \quad (8)$$

- rewiring a link with probability P_{decay} when agent i forwarded a query to a neighbor j in the stage $t - 1$, it rejects forwarding the query at stage t , and the number of cooperative neighbors is under a threshold σ .

$$a_i^t = \lambda \text{ if } a_i^{t-1} = j \wedge a_j^t = \emptyset \wedge |coop| < \sigma, coop \subseteq N_i(g) \quad (9)$$


4 Experiments

In this section we evaluate the effects of different criteria for the distribution of incentives and the social plasticity in the emergence of cooperation in a decentralized service discovery system.


The tests were performed on a set of 100 undirected networks based on preferences where the degree of connection followed an exponential distribution. The networks were populated by 1,000 agents and the average degree of connection was 2.5. Each agent had an initial budget $b = 100$. The agents played one role and offered one semantic web service related to this role. Initially, agents were uniformly distributed over 16 roles, which were defined in an organizational ontology. The set of semantic service descriptions used for the experiments was taken from the OWL-S TC4 test collection¹.

All the agents in the system had the same probability of generating service queries. A query was successfully solved when an agent that offered a similar service (i.e., the degree of semantic match between the semantic service descriptions and roles was over a threshold $\varepsilon = 0.75$) was found before the TTL ($TTL = 100$). Different values for


¹ <http://www.semwebcentral.org/projects/owl-s-tc/>




(a) Average budget per agent.




(b) Average budget per agent.




(c) Degree of collaboration.




(d) Degree of collaboration.




(e) % of successful searches.



(f) % of successful searches.



(g) Average path length.



(h) Average path length.

Fig. 2: Evolution of metrics when agents use incentives to promote cooperation in the system. Two initial scenarios are considered: (Left column) the initial cooperator agents were the 60% of the population, (Right column) the initial cooperator agents were the 40% of the population. The set of metrics considered were: the average budget per agent with an specific degree of connection, the degree of cooperation, the percentage of discovery processes that end before TTL, and the average number of steps in successful discovery processes.

ε and TTL were considered in previous experiments [23]. $\varepsilon = 0.75$ offers a balance between the degree of similarity and, the success rate and the average number of steps in the service discovery process. $TTL = 100$ offers the opportunity to different search configurations to improve their success rate.

The query distribution in the system was modeled as a uniform distribution. In the experiments, we made a snapshot of all of the metrics every time 5,000 queries were solved in the system in order to see their evolution. In all the experiments we did 20 snapshots. The costs, benefits, and incentives of the actions were the following: $\beta = 0.5$, $p = 0.5$, $c = 0.01$, $\alpha = 0.02$ (when incentives are distributed uniformly), and $\gamma = 0.1$. For the mechanisms that distribute the incentives in a non-uniform way, agents distribute the quantity of 0.5 among the agents that participate in the discovery process considering the different criteria.

The metrics that we considered in the experiments were:

- the success of the service discovery process
- the path length of the discovery process
- the degree of cooperation in the system
- the budget that an agent has.

For the experiments we considered scenarios where incentives are used to promote cooperation and scenarios where the combination of incentives and social plasticity are used. Two initial configurations for the network population were also considered. In one configuration the 60% of agents were initially cooperators and in the other one the 40% of agents were cooperators.

4.1 Incentives

In these tests, we evaluated the different ways that an agent distributes the incentives among the agents that participated in the process when the percentage of initial cooperators changes. Figure 2 shows the results in two scenarios with different initial configurations. The left column shows the results from the scenario where the initial percentage of cooperator agents was 60%. The right column shows the results where the initial population of cooperators was the 40%. Figures 2a and 2b show the final budget of agents with certain degree after the last snapshot. The x-axis shows the degree of connection of the agents and y-axis shows the average budget that agents with certain degree of connection had available in the last snapshot. In general, agents with a high degree of connection were the agents that obtained higher benefits due to they participated in more service discovery processes, and usually, these processes were shorter and had more probability of success. The strategies that best distributed the incentives were the *Fixed* and the *InvSimDg* since it gave more incentive to those agents that were far from the target and had a low degree of connection. In Figure 2b, the budget for the nodes with a high degree of connection is higher than the results in Figure 2a, since in the majority of successful searches in the scenario where the initial percentage of cooperators was 40% the highly connected nodes participated.

Regarding the results related to the degree of cooperation in the system, the strategies that gave a fixed incentive to the participants in the discovery process obtained a

lower degree of cooperation than the strategies that did not distribute the incentives uniformly. In Figures 2c and 2d these results are shown. The x-axis shows the snapshots and the y-axis the number of agents that cooperate. The strategies that do not distribute uniformly the incentives benefit the highly connected agents of the network, this fact provides a higher degree of cooperation. In scenarios where the initial cooperators is high (60%), the strategy that uses fixed incentives also obtains good results. However, when the number of initial cooperators decreases, this strategy is not enough to promote cooperation in the system. As consequence of the cooperation promotion, the average number of steps required in the search process to reach the target agent decreases with strategies that do not distribute the incentives uniformly (see Figures 2e and 2f). Similarly, in scenarios where the initial cooperators are the 60% or the 40% of the initial population, the percentage of queries successfully solved increases with strategies that do not distribute uniformly the incentives (see Figures 2h and 2g).

4.2 Incentives and Social Plasticity

In these tests, we incorporated social plasticity. Agents used incentives to promote cooperation but also they rewired links that they considered that were not being useful. The values of the parameters in the decay function 4 are: 1 for the slope y and 7 for the displacement d .

The first set of experiments evaluates the influence of the parameter σ that regulates the social plasticity action based on the local information about the neighbors of the agent i . Specifically, sigma is the threshold of collaborative agents that an agent should have in order to consider the social plasticity action. We have considered the following values for σ : 0.25, 0.5, 0.75, and 0.95. For instance, the value $\sigma = 0.25$ means that the agent will consider the use of social plasticity mechanism if the number of collaborative agents in its neighborhood is under the 25%. Low values of σ means that the agent is going to use social plasticity only when the majority of its neighbors does not cooperate. High values of σ means that the agent is going to consider the use social plasticity although the majority of the neighbors are collaborative. We considered critical scenarios where the initial percentage of cooperators was the 40% of the network population.

Figure 3 shows the results obtained with different values of σ . With low values of σ , agents try to avoid the use of social plasticity and only use it in extreme scenarios where all the neighbors are not cooperative. For the incentive strategies that use a dynamic distribution of incentives the value of σ does not have a meaningful influence. However, in the strategies *Fixed* and *System*, a low value of σ ($\sigma = 0.25$) improves the cooperation degree of the network, but this improvement is not as significant as with higher values of σ such as 0.5 or 0.75. The best results are obtained with $\sigma = 0.75$. With $\sigma = 0.95$ agents use social plasticity almost always, therefore, the majority of non-cooperative nodes are isolated and the main component of the network loses more nodes than in other configurations where social plasticity is not used so frequently. As a result of this experiment, for the following experiments we use $\sigma = 0.75$.

Figure 4 shows the results obtained when incentives and social plasticity were combined in scenarios where the initial cooperators where the 60% or the 40% of the population. In general, it can be observed that the use of social plasticity improves the results

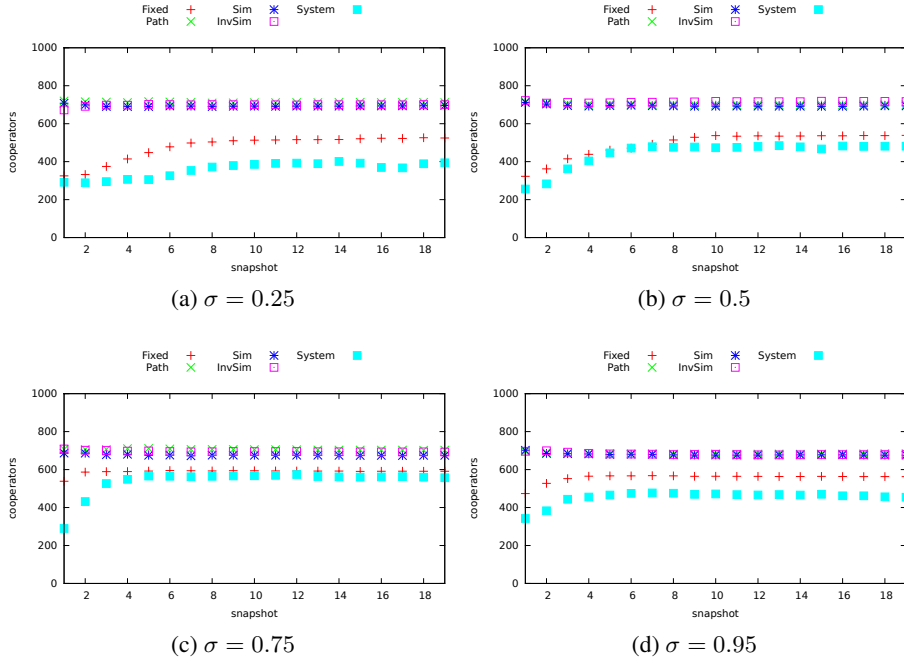




Fig. 3: Evolution of the degree of cooperation in the system when agents use incentives and have social plasticity with different values for the σ parameter.

obtained incentives where only considered and in scenarios where there was less percentage of initial cooperators. Regarding the final budget of the agents, the use of social plasticity implies a small decrease in the budget of the highly connected agents due to the number of cooperators with a lower degree of connection increases and highly connected nodes do not play a critical role in the search process.


The improvement obtained with the integration of social plasticity and incentives is more significant in the promotion of cooperation in scenarios where only the 40% of the population where initially cooperators. Specifically, the integration of social plasticity benefits the incentive strategies that use a uniform distribution of incentives (see Figures 4c and 4d). The increase of the degree of cooperation in the system facilitates the service discovery decreasing the average path length of the discovery processes (see Figures 4g and 4h) and increasing the success (see Figures 4e and 4f). This improvement is more significant than the improvement obtained only with the use of incentives. Finally, we analyzed the number of structural relations that were modified using different mechanisms for distributing the incentives (see Figures ?? and ??). The results show that the incentive mechanisms that distribute the benefit in a non uniform way require less structural changes to increase the cooperation in the system.




(a) Average budget per agent.




(b) Average budget per agent.




(c) Degree of collaboration.




(d) Degree of collaboration.




(e) % of successful searches.




(f) % of successful searches.




(g) Average path length.



(h) Average path length.



(i) Rewired structural relations.



(j) Rewired structural relations.

Fig. 4: Evolution of metrics when agents use incentives and social plasticity to promote cooperation in the system. Two scenarios are considered: (Left column) 60% of the initial population cooperate, (Right column) 40% of the initial population cooperate. The metrics considered were: the average budget per agent with an specific degree of connection, the degree of cooperation, % of searches that end before TTL, the average number of steps in successful searches, and the number of structural relations modified.

5 Conclusions

This article addresses the problem of emergence of cooperation in scenarios where cooperation is required to achieve a good performance that benefits all of the participants. Specifically, our proposal focuses on the emergence of cooperation in decentralized service discovery scenarios where agents need the cooperation of their neighbors in order to locate other agents that offer services that they require. Therefore, if selfish agents appear in the system, in the long term, as the number of non-cooperator agents increases, the service discovery process could be seriously compromised. For this reason, it is important to provide mechanisms that facilitate the emergence and maintenance of cooperation. In this paper, we present the combination of two mechanisms to facilitate the emergence of cooperation in open societies of agents where there are cooperative and non cooperative agents and they can change their behavior.

In the model that we presented, agents can use incentives in order to promote cooperative actions such as the forwarding action in the discovery process. We have considered different mechanisms to distribute these incentives. Some of them take into account the same quantity of reward for all the participants in a successful search process. Others distribute the reward among participants non-uniformly considering an specific criterion. The experiments shown that the strategies *InvSim* and *Fixed* distribute the benefits more uniformly among all the agents than the others. Moreover, the strategies that promote a higher degree of cooperation in the network are non-uniform.

Moreover, we also considered the inclusion of structural changes (social plasticity) based on the degree of cooperation of their neighbors. As the number of times a neighbor refuses to forward a query increases, the probability of changing this relation increases. If an agent decides to change a neighbor, it chooses a neighbor with similar functional features to the previous one. The inclusion of social plasticity in the system increases the degree of cooperation achieved in the system, mainly when the incentive mechanism used is based on a fixed reward distribution. Besides that, the use of social plasticity decreases the differences among uniform and non-uniform strategies.

The experiments confirm that this combination of mechanisms promote cooperation in different service discovery scenarios where the number of cooperators is higher than the number of non-cooperators or where the number of non-cooperator agents is higher than the number of cooperators. The increase of the degree of cooperation in the system improves the performance of the system reducing the average number of steps required to reach the target and increasing the number of service discovery processes. As future work, we are going to consider new distribution strategies for incentives. Moreover, we are thinking in distributing the cost of social plasticity action among the agents that have a connection with a non-cooperator agent and how this distribution affects to the service discovery performance. Finally, we would like to evaluate other types of network structures and analyze how this influences in the cooperation emergence.

6 Acknowledgements

Work partially supported by the Spanish Ministry of Science and Innovation through grants TIN2009-13839-C03-01, TIN2012-36586-C03-01, CSD2007-0022 (CONSOLIDER-INGENIO 2010).

References

1. Sun, Q., Garcia-Molina, H.: Slic: A selfish link-based incentive mechanism for unstructured peer-to-peer networks. In: Proceedings of the 24th International Conference on Distributed Computing Systems (ICDCS'04), Washington, DC, USA, IEEE Computer Society (2004) 506–515
2. Lin, W., Zhao, H., Liu, K.: Incentive cooperation strategies for peer-to-peer live multimedia streaming social networks. *IEEE Transactions on Multimedia* **11**(3) (april 2009) 396–412
3. Gu, B., Jarvenpaa, S.: Are contributions to p2p technical forums private or public goods? - an empirical investigation. In: Proceedings of the 1st Workshop on Economics of Peer-to-Peer Systems. (2003)
4. Shneidman, J., Parkes, D.C.: Rationality and self-interest in peer to peer networks. In: Proceedings of the 2nd Int. Workshop on Peer-to-Peer Systems (IPTPS'03). (2003)
5. Blanc, A., Liu, Y.K., Vahdat, A.: Designing incentives for peer-to-peer routing. In: Proceedings of the 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Volume 1. (march 2005) 374–385 vol. 1
6. Hardin, G.: The tragedy of the commons. *Science* (162) (1968) 1243–1248
7. Doran, J.E., Franklin, S., Jennings, N.R., Norman, T.J.: On cooperation in multi-agent systems. *The Knowledge Engineering Review* **12** (1997) 309–314
8. Nowak, M.A.: Five Rules for the Evolution of Cooperation. *Science* **314**(5805) (2006) 1560–1563
9. Nowak, M.A., Sigmund, K.: Evolution of indirect reciprocity by image scoring. *Nature* **393**(6685) (1998) 573–577
10. Sigmund, K.: Sympathy and similarity: The evolutionary dynamics of cooperation. *Proceedings of the National Academy of Sciences* **106**(21) (2009) 8405–8406
11. Hauert, C., Traulsen, A., Brandt, H., Nowak, M.A., Sigmund, K.: Via Freedom to Coercion: The Emergence of Costly Punishment. *Science* **316**(5833) (2007) 1905–1907
12. Sigmund, Hauert, C., Nowak, M.: Reward and punishment. *P Natl Acad Sci USA* (19) (2001) 10757–10762
13. Sigmund, K.: Punish or perish? retaliation and collaboration among humans. *Trends in Ecology and Evolution* **22**(11) (2007) 593–600
14. Pujol, J.M., Delgado, J., Sangüesa, R., Flache, A.: The role of clustering on the emergence of efficient social conventions. In: *IJCAI*. (2005) 965–970
15. Ohtsuki, H., Hauert, C., Lieberman, E., Nowak, M.A.: A simple rule for the evolution of cooperation on graphs and social networks. *Nature* **441**(7092) (2006) 502–505
16. Santos, F.C., Santos, M.D., Pacheco, J.M.: Social diversity promotes the emergence of cooperation in public goods games. *Nature* **454**(7201) (2008) 213–216
17. Hofmann, L.M., Chakraborty, N., Sycara, K.: The evolution of cooperation in self-interested agent societies: a critical study. In: Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2. (2011) 685–692
18. Eguluz, V.M., Zimmermann, M.G., Cela-Conde, C.J., Miguel, M.S.: Cooperation and emergence of role differentiation in the dynamics of social networks. *American Journal of Sociology* **110** (2005) 977
19. Axelrod, R.M.: The evolution of cooperation. Basic Books, New York (1984)
20. Griffiths, N., Luck, M.: Changing neighbours: improving tag-based cooperation. In: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1 - Volume 1. (2010) 249–256
21. Villatoro, D., Sabater-Mir, J., Sen, S.: Social instruments for robust convention emergence. In Walsh, T., ed.: Proceedings of the International Joint Conference on Artificial Intelligence. (2011) 420–425

22. Del Val, E., Rebollo, M., Botti, V.: Enhancing Decentralized Service Discovery in Open Service-Oriented Multi-Agent Systems. *Journal of Autonomous Agents and Multi-Agent Systems* (2012)
23. del Val, E.: *Semantic Service Management in Service-Oriented Multi-Agent Systems*. PhD thesis, Departament de Sistemes Informàtics i Computació, Universitat Politècnica de València (2013)