Discussion Papers in Economic Behaviour

DPEB
02/14

Analysis of Strategies to Promote Cooperation in Distributed Service Discovery

>Guillem Martínez-Cànovas
Universitat de València and ERI-CES, Spain
>Elena Del Val
DSIC, Universitat Politècnica de València, Spain
>Vicente Botti
DSIC, Universitat Politècnica de València, Spain
>Penelope Hernández
Universitat de València and ERI-CES, Spain
>Miguel Rebollo
DSIC, Universitat Politècnica de València, Spain

March, 2014

www.erices.es

Analysis of Strategies to Promote Cooperation in Distributed Service Discovery

Guillem Martínez-Cánovas[b], Elena Del Val[a], Vicente Botti[a], Penélope Hernández[b], Miguel Rebollo[a]

[a]*Departament de Sistemes Informàtics i Computació, Universitat Politècnica de València, Spain*
[b]*ERI-CES, Departamento de Análisis Económico, Universitat de València, Spain*

Abstract

New systems can be designed, developed, and managed as societies of agents that interact with each other by offering and providing services. These systems can be viewed as complex networks where nodes are bounded rational agents. In order to deal with complex goals, they require the cooperation of the other agents to be able to locate the required services. In this paper, we present a theoretical model that formalizes the interactions among agents in a search process. We present a repeated game model where the actions that are involved in the search process have an associated cost. Also, if the task arrives to an agent that can perform it, there is a reward for agents that collaborated by forwarding queries. We propose a strategy that is based on random-walks, and we study under what conditions the strategy is a Nash Equilibrium. We performed several experiments in order to validate the model and the strategy and to analyze which network structures are more appropriate to promote cooperation.

Keywords: Repeated games, Networks, Nash Equilibrium, Random-walk, Service Discovery

---

Email addresses: `guillem.martinez@uv.es` (Guillem Martínez-Cánovas), `edelval@dsic.upv.es` (Elena Del Val), `vbotti@dsic.upv.es` (Vicente Botti), `penelope.hernandez@uv.es` (Penélope Hernández), `mrebollo@dsic.upv.es` (Miguel Rebollo)

Preprint submitted to Nuclear Physics B *March 22, 2014*

1. Introduction

Social computing has emerged as a discipline in different fields such as Economics, Psychology, and Computer Science. Computing can be seen as a social activity rather than as an individual one. New systems are designed, developed, and managed as societies of independent entities or agents that offer services and interact with each other by providing and consuming these services [18]. These systems and applications can be formally represented through formal models from the field of Complex Networks [22]. This area provides a sound theoretical basis for the development of models that help us to reason about how distributed systems are organized [14]. Complex Network models have been used in different contexts such as social networks (collaboration, music, religious networks), economic networks (trade, tourism, employment networks), Internet (structure and traffic networks), bio-molecular networks, and computer science networks among others [5, 21].

In systems of this kind, one of the challenges is the design of efficient search strategies to be able to locate the resources or services required by entities in order to deal with complex goals [2, 21, 7]. Taking into account the autonomy of the entities that participate in the search process, three levels of search decentralization can be considered. We consider that at the first level the search process is centralized when there is a common protocol that is adopted by all the entities of the system and this protocol dictates the actions that must be followed (i.e., the protocol specifies the entity that starts the process, the sequence of participation of entities, and the target). At the second level this protocol can be relaxed. The entities adopt that protocol, and, therefore, they carry out the same set of actions, but the search path (i.e., the sequence of entities that participate in the search process) is not specified. At the third level, a decentralized search can be considered when there is a protocol adopted by all the entities that specifies the set of available actions. However, these entities can decide whether or not they are going to follow the protocol. It would not be desirable that to impose the same behavior on all the nodes if it takes away their individual choice, (i.e., it would be desirable that all the nodes would follow the protocol willingly). Therefore, we have looked for a concept of stability within the strategies of the entities of the system. This concept, which comes from Game Theory, is known as Nash Equilibrium.

As an application scenario, we consider a P2P system that is modeled as a multi-agent system. Agents act on behalf of users playing the role of a service provider or service consumer. Agents that play the role of service consumers should be able to locate services, make contracts agreements, and receive and present results [28]. Agents that play the role of service providers should be able to manage the access to services and ensure that contracts are fulfilled. By considering the system as a network, it is assumed that all the information is distributed among the agents. Since agents only have a local view of the network, the collaboration of other agents is required in order to reach the target. During a search process, agents can carry out a set of actions: create a task that must be performed by a qualified agent, forward the task to one or several neighbors if they do not know how to solve the task, or perform the task if they can provide the required service. The cooperation of agents forwarding queries plays a critical role in the success of the search process [6]. This action facilitates the location of a resource based on local knowledge. However, in our scenario, this action has an associated cost and agents are free to decide whether or not the forwarding action is profitable to them based on its cost and the expected reward.

In this paper, we propose a model to formally describe the distributed search for services in a network as a game. Specifically, we use the repeated games framework to model both the process that a task follows through the network and the global task-solving process. In the former, each period is a decision stage for the agent who is in possession of the task. In the latter, a project is generated in each period and randomly assigned to an agent in the network.

Our intention is to analyze the relationship between the cost of forwarding the task and the reward that agents obtain later when the task is solved in order to guarantee that cooperation is a stable behavior in the game. We called this reward α. We establish a bound for the total length of the search process total length using Mean First Passage Time (MFPT), which is the average number of steps necessary to go from an agent i to another agent j in the same network. Therefore, the structure of the network also characterizes α through the MFPT, and, consequently, the network structure influences the agents' behavior. In order to verify this, we ran simulations to contrast the possible differences among network structures. The results show that the structure of the network has a significant influence on the emergence of cooperation. The structure that offers the best results is the Scale-Free structure since its diameter is closer to the limit of steps in the search process than the other network structures.

The paper is organized as follows. Section 2 presents a repeated game model to formalize the search process of services in agent networks. In Section 3, some strategies that agents can follow in the repeated game are analyzed in order to determine whether or not they are at a Nash Equilibrium. Section 4 describes several experiments we performed to empirically validate the theoretical results in different network structures as well as to analyze the influence of the network structure and to determine which structure facilitates the emergence of cooperation in the proposed repeated game. Section 5 presents other works related to cooperation emergence in distributed environments. Finally, Section 6 presents the conclusions.

2. The Model

Consider a finite set of agents $N = \{1, 2, \ldots, n\}$ that are connected by undirected links in a fixed network represented by the adjacency matrix g. A link between two agents i and j, such that $i, j \in N$, 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

$$N_i = \{j | g_{ij} = 1\}$$

For simplicity we assume that $g_{ii} = 0$ so all neighbors in $N_i(g)$ are different from i. The number of neighbors that agent i has (its degree of connection) is denoted by $k_i = |N_i(g)|$, which is the cardinality of the set $N_i(g)$. Alternatively, we use the adjacency matrix to represent the network, which is denoted by A. A link between agents i and j is represented by $A_{ij} = 1$, and by $A_{ij} = 0$ if there is no link.

We consider that each agent has a type (service) $\theta_i \in [0, 1]$ that represents the degree of ability of agent i. Let $\rho \in [0, 1]$ be a task that must be carried out by one of the agents in the network. We assume that there is at least one agent $i \in N$ such that i is suitable to perform the task, which means that, for a fixed ε, its type θ_i is 'similar' to the task ρ, *i.e.*, $|\theta_i - \rho| \leq \varepsilon$.

We define an N-person network game Γ_ρ^∞ that takes place in g. Each agent has a set of actions $A_i = \{\emptyset, 1, 2, \ldots, N_i, \infty\}$, where:

- ∞ means the agent itself does the task
- $\{1, 2, \ldots, N_i\}$ means forwarding the task to one of the agent's N_i neighbors
- $\emptyset$ means doing nothing

In the first period of the game, a task ρ is uniformly assigned to a randomly selected agent. Beginning at stage 1, the task passes through the network stage by stage. At stage $t > 0$ each agent chooses one of the above actions depending on whether or not the task is in the agent's node. The action perform by an agent where the task is not in the agent's node is considered to be $\emptyset$ or "doing nothing". We associate a null payoff to this action. At stage 1, agent $i(1)$ chooses one action from its action set.

At the initial stage, if the first agent, $j \in \{1, \ldots, N\}$ chooses to do the task itself because its type is ε close to the task ρ, then the game ends. Agent j gets a payoff of $1 - |\rho - \theta_j|$, which depends on its type θ_j and the task ρ. The more similar the type and the task are, the greater the payoff is. The rest of the agents can do any action in their action sets. More specifically, let $c > 0$ be the cost of forwarding the task. If an agent forwards the task, at some point it, the agent may earn a payoff $\alpha > c$ if the task ends successfully. If an agent chooses the action $\emptyset$, the payoff is 0 if the agent did not forward any task in a previous period or the agent did forward the task but the task ended unsuccessfully (i.e.,nobody chose ∞).

Formally:

$$u_i^t(a_i, a_{-i}; j) = \begin{cases} 1 - |\rho - \theta_i| & \text{if } a_i^t = \infty \\ -c & \text{if } a_i^t \in \{1, \ldots, N_i\} \\ 0 & \text{if } a_i^t = \emptyset \wedge \nexists t' < t : a_i^{t'} \in \{1, \ldots, N_i\} \\ \alpha & \text{if } a_i^t = \emptyset \wedge \exists t' < t : a_i^{t'} \in \{1, \ldots, N_i\} \wedge \exists j \in N : a_j^t = \infty \end{cases}$$

By choosing actions at stage t, agents are informed of actions that are chosen in previous stages of the game. Therefore, let us consider a complete information set-up. Formally, let $H_t, t = 1, \ldots,$ be the cartesian product $A \times A$ $t-1$ times, i.e., $H_t = A^{t-1}$, with the common set-theoretic identification $A^0 = \oslash$, and let $H = \cup_{t \geq 0} H_t^{\infty}$. A pure strategy σ^i for agent i is a mapping from H to $A^i, \sigma^i : H \to A^i$. Obviously, H is a disjoint union of H_t, $t = 1, \ldots, T$ and σ_t^i:$H_t \to A^i$ as the restriction of σ^i to H_t.

The payoff function of each agent when the game is repeated a certain number of times and when the task starts at any agent is formalized as:

$$u_i(\sigma_i, \sigma_{-i}) = \sum_{t>0} u_i^t(a_i, a_{-i}; j) = \frac{1}{N} \sum_{j=1}^{N} \sum_{t=1}^{\infty} u_i^t(a_i, a_{-i}; j)$$

This induces an order in the payoffs for each strategy σ_i that each agent

i chooses given the action profiles σ_{-i}, which allows us to rank them and calculate the Nash Equilibriums of the game.

In order to characterize the set of feasible and individual rational level to define the set of equilibria payoffs of the repeated game (*i.e.*, the equilibrium payoff attained as a consequence of the well-known Folk Theorem) [9], we have to establish the min-max level in pure actions. The min-max strategy for agent i is the one that guarantees the highest possible payoff in the action profile that is the worst case scenario for agent i. This is sometimes called the reservation payoff. Formally,

$$\bar{u_i} = \min_{a_{-i}} \max_{a_i} u_i(a_i, a_{-i}), a_i \in A_i, a_{-i} \in A_{-i} \tag{1}$$

In our set-up for the one-shot payoff function, the min-max strategy is $\emptyset$ and, therefore, the min-max level is 0. An action profile $(\sigma_1^*, \dots, \sigma_n^*)$ is a Nash equilibrium in the network game Γ, if and only if

$$u_i(\theta_i, \sigma_1^*, \dots, \sigma_n^*) \geq u_i(\theta_i, \sigma_1^*, \dots, \hat{\sigma}_i, \dots, \sigma_n^*), \quad \forall\, \sigma_i^* \neq \hat{\sigma}_i,\ i \in N,\ \text{and}\ \theta_i \in \Theta.$$

We define the set of feasible payoff vectors as

$$F := \operatorname{conv}\{u(a), a \in \mathcal{A}\}.$$

The set of *strictly individually rational payoff vectors* (relative to the min-max value in pure strategies) is

$$V := \{x = (x_1, \dots, x_n) \in F \colon x_i > \bar{u}_i \ \ \forall i \in N\}.$$

Folk theorems in the context of game theory establish feasible payoffs for repeated games. Each Folk Theorem considers a subclass of games and identifies a set of payoffs that are feasible under an equilibrium strategy profile. Since there are many possible subclasses of games and several concepts of equilibrium, there are many Folk Theorems.

The Folk Theorem states any payoff profile in V can be implemented as a Nash equilibrium payoff if δ is large enough. The intuition behind the Folk Theorem is that any combination of payoffs such that each agent gets at least its min-max payoff is sustainable in a repeated game, provided each agent believes the game will be repeated with high probability. For instance, the punishment imposed on an agent who deviates is that the agent will be held to its min-max payoff for all subsequent rounds of the game. Therefore, the short-term gain obtained by deviating is offset by the loss of payoff in future rounds. Of course, there may be other, less radical (less grim) strategies that also lead to the feasibility of some of those payoffs. The good news from the Folk Theorem is that a wide range of payoffs may be sustainable in equilibrium. The bad news is that, there may exist a multiple number of equilibria.

3. Equilibrium strategies

In this section, we study which strategy profiles are a Nash Equilibrium in the game Γ_ρ^∞. Namely, we start defining the *Nobody works* strategy, which basically consists of doing nothing, even in the case that an agent can perform the task. We prove that the *Nobody works* strategy is not a Nash Equilibrium in the game. Then we consider the so-called random-walk strategy. In this strategy, an agent is not able to solve the project, it uniformly and randomly chooses one of its neighbors to forward the task to. We establish the conditions under which the strategy profile every agent plays the random-walk strategy is a Nash Equilibrium. We enrich the model by adding a threshold for the number of times that a task can be forwarded and we also study under which conditions is a Nash Equilibrium.

3.1. Nobody works

One possible strategy is the strategy we call *Nobody works*, in which every agent always chooses the action $\emptyset$ and consequently gets a payoff of 0. One of our model's assumptions is that for all possible task ρ there exists an agent that is able to perform it. Let that agent be i, and let its type be θ_i. From our payoff criterion, we can state that in some period t the project will start at agent i and agent i will be able to solve it. In that case, if agent i chooses the ∞ action (doing the project), agent i gets a payoff of $1 - |\rho - \theta_i| > 0$; therefore, the *Nobody works* strategy is not an equilibrium strategy.

3.2. Random Walk

In this subsection, we study the case where all agents play a behavioral strategy $\sigma_i^\infty : H^{t-1} \to \Delta(A_i)$, which leads to the well-known dynamics of "random-walk". We call this behavioral strategy the random-walk strategy.

Let us formally define the random-walk strategy. At each stage t, agent i performs one of the three actions that are possible:

- the $\emptyset$ action if no task arrives.

- the ∞ action if agent i's type θ_i is close to the task ρ_t.
- the forwarding action when the task arrives and agent i cannot solve it, agent i uniformly and randomly chooses one of its neighbors to forward the task to.

This strategy is a "myopic" strategy since agents do not update the expected payoff. Each agent i will uniformly and randomly choose one of its neighbors to continue searching for the agent that can solve the task ρ. Recall that in our game for all task ρ, there exists an agent k^* such that agent i can do the task ρ (*i.e.*, $|\theta_{k*} - \rho| > \varepsilon$). As a consequence of the random-walk strategy we can assert the existence of a finite time $0 \leq \tilde{t} < \infty$ and $k^* \in \{1, \ldots, N\}$ such that $a^{\tilde{t}}_{K*} = \infty$. Therefore, given a task ρ, the achieved payoff for each agent first depends on whether or not agent i was part of the path of searching for the agent that did the task. If agent i did not in the procedure, then agent i gets 0, which is the $\min - \max$ value.

Now, suppose that i is part of the path. Let us define some parameters that take part in the utility function. We refer to the probability of an agent being capable of performing the task as γ_i, and since it is the same for all agents , we simply call it simply γ. P^∞_x is the probability that the task reaches a specific agent x in the long run, and the previously defined parameters α and c are the reward and the cost of forwarding the task, respectively.

Hence, the utility function of the game Γ^∞_ρ for agent i is:

$$u_i(\theta_i, \sigma_i, \sigma_{-i}) = P^\infty_i \left(\gamma(1 - |\rho - \theta_i|) + (1 - \gamma)(P^\infty_{k^*}(\alpha - c) + (1 - P^\infty_{k^*})(-c)\right)) \tag{2}$$

The following proposition states that the strategy profile in which every agent plays a random-walk strategy is a Nash equilibrium in the game Γ^∞_ρ.

Proposition 1. *The strategy profile* $(\sigma^\infty_1, \ldots, \sigma^\infty_n)$ *is a Nash Equilibrium in the game* Γ^∞_ρ.

Proof. Let i be an agent such that agent i selects a strategy $\sigma^\emptyset_i \neq \sigma^\infty_i$; let t be a time period such that the task ρ arrives to i; let t' be another time period such that $t' \neq t$ and let $|\theta_i - \rho| < \varepsilon$. The strategy $\sigma^\emptyset_i$ is formally defined as

$$(\sigma_i^{\emptyset})^t : H^t \to A_i \begin{cases} (\sigma_i^{\emptyset})^t & = \infty \\ \forall t \neq t', (\sigma_i^{\emptyset})^t & = \emptyset \end{cases} \tag{3}$$

When selecting that strategy, if i is able to afford the task agent i does it, and that is the only profit that agent i eventually gets because it never forwards the task. Consequently, agent i's utility function is

$$u_i(\theta_i, \sigma_i^{\emptyset}, \sigma_{-i}^{\infty}) = P_i^{\infty}\left(\gamma(1 - (\rho - \theta_i))\right) \tag{4}$$

In order to prove that the strategy profile $(\sigma_1^{\infty}, \ldots, \sigma_n^{\infty})$ is a Nash Equilibrium the utility function described in 2 must be greater or equal to the utility function specified in 4

$$\begin{aligned} P_i^{\infty}\left(\gamma(1 - (\rho - \theta_i)\right) + (1-\gamma)(P_{k*}^{\infty}(\alpha - c) + (1 - P_{k*}^{\infty})(-c)) &\geq P_i^{\infty}\left(\gamma(1 - (\rho - \theta_i))\right) \\ (1-\gamma)(P_{k*}^{\infty}(\alpha - c) + (1 - P_{k*}^{\infty})(-c)) &\geq 0 \end{aligned}$$

Since $0 < \gamma < 1$, $(1-\gamma)$ is always positive. Then

$$\begin{aligned} P_{k*}^{\infty}(\alpha - c) + (1 - P_{k*}^{\infty})(-c) &\geq 0 \\ \alpha &\geq \frac{c}{P_{k*}^{\infty}} \end{aligned}$$

By the definition of random-walk dynamics, in the long term, a task ρ will always find the agent k^* that is capable of solving it, so $P_{k*}^{\infty} = 1$. Since $\alpha > c$ by assumption, the strategy profile $(\sigma_1^{\infty}, \ldots, \sigma_n^{\infty})$ is a Nash Equilibrium in the game Γ_{ρ}^{∞}.

□

Now we enrich the model by introducing a "time" condition to solve the task. It makes sense to limit the rewards for efforts to solve or forward the task to a time limit within which the task must be solved (i.e., efforts are only rewarded if the task is solved in a certain number of time periods).

3.3. Random-walk strategy with a finite number of steps

An interesting measure for establishing the limit of steps that a task ρ can take to be solved is the Mean First Passage Time (hereafter MFPT). The MFPT between two nodes i and j of a network is defined as the average number of steps to go from i to j in that particular network [32]. Therefore, we define the strategy σ_i^{τ} for an agent i, which consists of forwarding the task to a randomly selected neighbor only if it has advanced a number $t_i < \tau$ times, where τ is the average MFPT of the network (which we formally define below).

From equation 34 of [32], we know that the MFPT from any agent to a particular agent j in a network is defined as:

$$\langle T_j \rangle = \frac{1}{1-\pi_j} \sum_{i=1}^{N} \pi_i T_{ij} =$$
$$= \frac{1}{1-\pi_j} \sum_{k=2}^{N} \left(\frac{1}{1-\lambda_k} \psi_{kj}^2 \sum_{i=1}^{N} \frac{k_i}{k_j} \right) - \frac{1}{1-\pi_j} \sum_{k=2}^{N} \left(\frac{1}{1-\lambda_k} \psi_{kj} \sqrt{\frac{K}{k_j}} \sum_{i=1}^{N} \psi_{ki} \sqrt{\frac{k_i}{K}} \right)$$

where k_i and k_j are the degree of agents i and j, respectively, ψ_k is the kth eigenvector of S corresponding to the kth eigenvalue λ_k (with $S = D^{-\frac{1}{2}} A D^{-\frac{1}{2}}$, A being the adjacency matrix of the network, D being the diagonal degree matrix of the network, and the eigenvalues being rearranged as $1 = \lambda_1 > \lambda_2 \geq \lambda_3 \geq \ldots \geq \lambda_N \geq -1$) and $\pi_j = d_j/K$ (with $K = \sum_{j=1}^{N} d_j$).

It follows from Eq. (6) from the same article that $\sum_{i=1}^{N} \psi_{ki} \sqrt{\frac{k_i}{K}} = \sum_{i=1}^{N} \psi_{ki} \psi_{1i} = 0$. Thus, the second term is equal to 0. So

$$\langle T_j \rangle = \frac{1}{1-\pi_j} \sum_{k=2}^{N} \left(\frac{1}{1-\lambda_k} \psi_{kj}^2 \sum_{i=1}^{N} \frac{k_i}{k_j} \right) = \frac{1}{1-\pi_j} \frac{K}{k_j} \sum_{k=2}^{N} \frac{1}{1-\lambda_k} \psi_{kj}^2 \tag{5}$$

We define the maximum number of steps that must be taken for every task ρ to be solved as the average $\langle T_j \rangle$ for all $j \in N$, which we denote as τ. Formally:

$$\tau = \sum_{j=1}^{N} \frac{\langle T_j \rangle}{N} \tag{6}$$

We now define a new game Γ_ρ^τ. In this game, if it takes more than τ steps to solve the task, the game ends and the collaborating agents get no reward. In the following, we explain the equilibrium strategies for the game Γ_ρ^τ.

Let us define some new parameters that play a role in the new game: the number of steps a task has advanced until it reaches agent i is t_i; $Q_{i,k*}^{\tau - t_i}$ is the probability that the task reaches an agent k^* starting from agent i in

$\tau - t_i$ or less steps and $P^{\tau}_{s,i}$ is the probability that the task reaches an agent i starting from agent s in τ or less steps.

In order to formally define $P^{\tau}_{s,i}$, we use the adjacency matrix of the network (denoted as A) and one of its properties which states that the values (i, j) of A^n indicate the number of paths of length n between i and j in that network. $P^{\tau}_{s,i}$ can be defined as the number of paths of length τ or less between s and i divided by the total number of paths with the same length starting at s but ending at any possible agent j of the network. Formally:

$$P^{\tau}_{s,i} = \frac{\sum\limits_{t=1}^{\tau}(A^t)_{si}}{\sum\limits_{j=1}^{N}\left(\sum\limits_{t=1}^{\tau}(A^t)_{sj}\right)} \tag{7}$$

Let us define $r_i^{\tau-t_i}$, or simply r_i, as the number of agents that the task can reach starting from i in $\tau - t_i$ or less steps. For this purpose, we use the Reachability matrix, (denoted as R), which is defined as

$$\forall i, j \in N, (R^{\tau-t_i})_{ij} = \begin{cases} 1 & \text{if there exists at least one path between } i \text{ and } j \\ & \text{ of length } \tau - t_i \text{ or less} \\ 0 & \text{otherwise} \end{cases} \tag{8}$$

The process for obtaining R from the adjacency matrix is straightforward. Then, we formally define r_i as

$$r_i = \sum_{j=1}^{N}(R^{\tau-t_i})_{ij} \tag{9}$$

To define $Q^{\tau-t_i}_{i,k*}$, we compute the probability that none of the reachable agents for agent i is able to solve the task, which is $(1-\gamma)^{r_i}$. Then $Q^{\tau-t_i}_{i,k*}$ is defined as

$$Q^{\tau-t_i}_{i,k*} = 1 - (1-\gamma)^{r_i} \tag{10}$$

Hence, the utility function of the game Γ^{τ}_{ρ} for an agent i when all agents play the strategy σ^{τ} is

$$u_i(\theta_i, \sigma_i^\tau, \sigma_{-i}^\tau) = P_{s,i}^\tau \left(\gamma(1-(\rho-\theta_i)) + (1-\gamma)(Q_{i,k*}^{\tau-t_i}(\alpha-c) + (1-Q_{i,k*}^{\tau-t_i})(-c)\right) \tag{11}$$

Now we study a bound for α for which the strategy profile $(\sigma_1^\tau, \ldots, \sigma_n^\tau)$ is a Nash Equilibrium in the game Γ_ρ^τ.

Proposition 2. *If* $\alpha_i \geq \dfrac{c}{1-(1-\gamma)^{r_i}}$, *the strategy profile* $(\sigma_1^\tau, \ldots, \sigma_n^\tau)$ *is a Nash Equilibrium in the game* Γ_ρ^τ*.*

Proof. The proof proceeds exactly like the proof for Proposition 1 but substituting the proper probabilities $P_{s,i}^\tau$ and $Q_{i,k*}^{\tau-t_i}$. Finally, we have

$$\alpha \geq \frac{c}{Q_{i,k*}^{\tau-t_i}} \tag{12}$$

By substituting 10 in 12, we have

$$\alpha \geq \frac{c}{1-(1-\gamma)^{r_i}} \tag{13}$$

The fact that α depends on r_i implies that each agent has its own bound for alpha which depends on the agent's connectivity (α becomes α_i).

□

This means that the network structure has a deep impact on α bounds. In high clustered networks, r_i is high for each agent i, and, consequently, α_i is low. The opposite occurs in low clustered networks (e.g., Erdös-Renyi networks) where α_i is uniform among all agents. In networks with a non-uniform degree distribution (e.g., scale-free networks), average α may be similar to the α for Erdös-Renyi networks, but it varies a lot between hub and terminal agents.

4. Experiments

In this section, we validate the proposed mathematical model for service search in different network structures. Specifically, we focus on how the structural parameters of the networks influence the required reward α in order to promote cooperation (i.e., forwarding tasks) and improve the success of the

topology	N	edges	avDg	std	clust	dens	τ = Log(mfpt)	d	diameter/τ
Random	100.0	200.00	4.00	1.46	0.02	0.04	5.00	7.09	1.418
	100.0	300.00	6.00	1.82	0.03	0.06	4.00	5.00	1.25
ScaleFree	100.0	197.00	3.94	3.94	0.02	0.04	5.00	5.09	1.09
	100.0	293.00	5.86	5.08	0.04	0.06	5.00	4.27	0.854
SmallWorld	100.0	200.00	4.00	1.02	0.08	0.04	5.00	7.55	1.55
	100.0	300.00	6.00	1.27	0.10	0.06	4.00	5.45	1.36

Table 1: Network structural properties: topology, number of agents, number of edges, average degree of connection of agents, standard deviation of the degree distribution, clustering, density, τ = Log(Mean First Passage Time),diameter, ratio diameter/τ.

search process. The structural parameters are represented by the parameter r_i (see Equation 9). For the evaluation, we compare the success rate of the searches and the average agent utility in different network structures. The network structures considered in the experiments are: Random, Scale-free, and Small-World networks.

4.1. Experimental Design

Each network in the experiments is undirected and has 100 agents. We also tested different sizes of networks, but the conclusions were similar to those obtained with 100 agents and we do not include them here. The structural properties of the networks are shown in Table 1. Each agent has a type (service) $\theta_i \in [0, 1]$ that represents the degree of ability of agent. The degree of ability is uniformly distributed among the agents. A task ρ is generated and assigned to an agent following a uniform probability distribution. Each agent has a set of actions to choose from when it receives a task: doing the task if the similarity between its ability and the task ρ is under a threshold $|\theta_i - \rho|$ ¡ ε, forwarding the task based on the expected reward (Formula 13), or doing nothing. The forwarding action has an associated cost $c = 5$. A task ρ is successfully solved when an agent that has an ability that is similar enough to the task ($|\theta_i - \rho| < \varepsilon$) in less than τ steps. For the experiments, the value of the τ is the Log(MFPT). We use this concave transformation to obtain clear results and illustrate the impact on the parameter with the structure of the network. The value for the ε parameter is 0.1. We executed each experiment over 10 networks of each type and we generated 1,000 tasks ρ in each network.

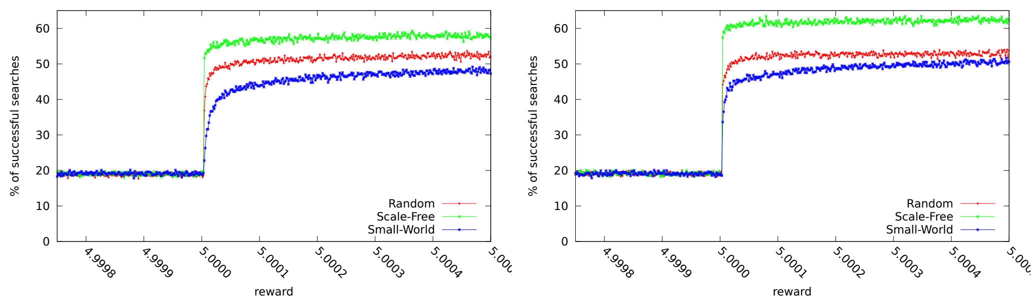

Figure 1: Influence of α values in the percentage of successful searches in different network structures of 100 agents. (Left) Network structures with an average degree of connection of 4. (Right) Network structures with an average degree of connection of 6.

4.2. The Influence of Structural Properties and α

In this section, we analyze the influence of network structural properties and reward α in the search process. We consider values for α in the range [4.99975, 5.0005] in order to see the effects on the search process (see Figure 1). In this interval, we observe the effects of considering values for α that are lower than the cost of the forwarding action (c=5), values that are equal to the cost of the forwarding action, and values that are greater than the cost of the forwarding action. With values of α lower than or equal to c, the success rate was around 20%. This percentage represents the number of tasks that can be solved directly by the first agent that receives the task. Values of α that are strictly superior to the cost of the forwarding action ($\alpha > c$) provide an increase in the success rate of the search process (see Figure 1 Left). The structural properties of the network considered in the search process have an important influence on the success rate. We observe that there are significant differences between the results in Scale-Free, Random, and Small-World networks. Scale-Free provided better results than the other networks since its structural properties increased the number of agents that could be reachable. The diameter of the network is closer to τ than the diameters of other network models (see Table 1). Another example of the influence of structural properties is the average degree of connection of the agents. As the average degree of connection increases, the number of reachable agents increases and so does the probability of finding the required agent. Therefore, agents estimate that it is profitable to forward the task to their neighbors (see Figure 1 Right).

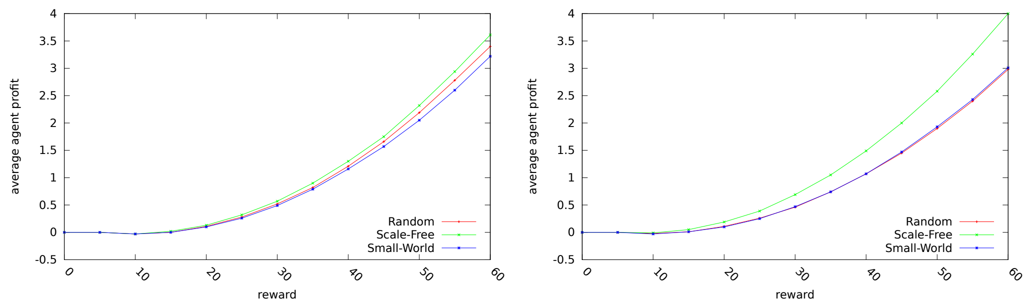

Figure 2: Influence of α values on the utility in different network structures of 100 agents. (Left) Network structures with an average degree of connection of 4. (Right) Network structures with an average degree of connection of 6.

The structural properties and the reward value α also influence the average utility obtained by an agent. In this experiment, we analyzed values of α in the range [0, 60]. We considered a wider range in order to see the values that made the average agent utility positive and how this utility evolves (see Figure 2). Values of α lower than or equal to c provided a utility equal to 0 since agents estimate that the expected reward was not enough to compensate the cost of the forwarding action. Values of α that were in the interval (5, 10] made some agents estimate that the forwarding action was going to be profitable. Although the value for α was enough for agents to consider forwarding tasks, their utility was not always positive for all the agents. Therefore, the average utility had a negative value. The interval (5, 10] for α values could be considered risky. The average utility became positive with α values greater than 10 (see Figure 2 Left). In this experiment, the network structure also had a significant influence. The Scale-Free network provided higher values of utility than the Random or Small-Word networks. This difference was also observed when we increased the average degree of connection of agents (see Figure 2 Right).

5. Related Work

Random-walk strategies have been presented as an alternative search strategy to *flooding* strategies [4, 31, 17] since they reduce the traffic in the system and provide better results [19, 33]. A random-walk search algorithm selects a neighbor randomly each time to forward the message to [10].

There are many search proposals that navigate networks using random-walk since they do no require specific knowledge and can be applied in several domains. Some of these works have introduced modifications such as using random-walk from multiples sources [34, 25] or adding information about routes [15, 1, 3] in attempt to improve the search efficiency. The influence of network structural properties on random-walk has also been studied. For instance, some of the properties that have been evaluated are: the mean first-passage time (MFPT) from one node to another [32, 27, 30], how the structural heterogeneity affects the nature of the diffusive and relaxation dynamics of the random walk [23], and the biased random-walk process based on preferential transition probability [8].

One of the common assumptions in network search is that all the agents have homogeneous behavior and that all of them are going to cooperate by forwarding messages. However, this does not correspond with real scenarios. In real large-scale networks, decisions are often made by each agent independently, based on that agent's preferences or objectives. Game Theoretic models are well suited to explain these scenarios [20]. Game theory studies the interaction of autonomous agents that make their own decisions while trying to optimize their goals. Game Theory provides a suite of tools that may be effectively used in modeling interactions among agents with different behaviors [29].

There are works in the area of Game Theory that focus on the routing problem in networks where there are selfish agents. Specifically, this problem has been studied in wireless and ad-hoc networks [29, 20]. Numerous approaches use reputation [13] (i.e., techniques based on monitoring the nodes' behavior from a cooperation perspective) or price-based techniques [12] (i.e., a node receives a payment for its cooperation in forwarding network messages and also pays other nodes which participate in forwarding its messages) to deal with selfish agents. One of the drawbacks of reputation systems is that nodes whose reputation values are higher than a threshold are treated equally. Therefore, a node can maintain its reputation value just above the threshold to obtain the same benefit as nodes with higher reputation levels. One of the problems of the Price-based techniques is that they are not fair with nodes located in region with low traffic that have few opportunities to earn credit. Li et al. [16] integrate both techniques and propose a game theory model for analyzing the integrated system. However, this approach does not consider the influence of the underlying structure in the cooperation emergence.

To understand the social behavior of the systems it is important to consider the network structure. There are several works that analyze the influence of the network structure when the agents of the networks do not follow homogeneous behavior. These works study how structural parameters such as clustering or degree distribution affect the emergence and maintenance of cooperative behavior among agents [26, 24]. Hofmann et al. [11] present a critical study about the evolution of cooperation in agent societies. The authors conclude that there is a dependence of cooperation on parameters such as network topology, interaction game, state update rules and initial fraction of cooperators.

The proposal presented in this paper analyzes through a game theory model the problem of cooperation emergence in the context of decentralized search. It differs from previous approaches in several ways. First, we considered a game that fits better with the characteristics of decentralized search than other games proposed in the literature that are based on the often studied Prisioner's Dilemma. Second, agents decision about cooperation is based on an utility function that takes into account the network topology properties. Moreover, the utility function also considers a limit in the number of possible steps to reach the target agent. This feature is important in distributed systems in order to avoid traffic overhead. Third, the strategy that agents follow is based on a search mechanism that is often used in network navigation and does not require specific domain knowledge. Therefore, the model can be easily applied in different search contexts. Finally, in order to promote cooperation, instead of using a reputation or price-based mechanisms, we use a mechanisms based on incentives provided by the system. We formally and experimentally analyze which is the minimum required reward in order to consider the strategy a Nash Equilibrium.

6. Conclusions

In this paper, we have analyzed the distributed search of resources in networks that model societies of agents. These agents offer services and interact with each other by providing and consuming these services. The actions of these agents have an associated cost and not all of the agents have homogeneous behavior. We have proposed the use of Game Theory to formally model the interactions between the agents as a repeated game, and we have described a strategy that is based on the simple well-known random-walk strategy. We have also established the conditions under which the random-walk strategy is a Nash Equilibrium. The strategy proposed has been extended by adding a constraint for contexts where the number of times a task can be forwarded is restricted. The conditions under which this extended strategy is a Nash Equilibrium have also been analyzed. Finally, we validated the proposed model and the latest strategy in different types of networks. The results show that in order to promote cooperation among the agents of the network, the expected reward should be greater than the cost of the forwarding action. Moreover, the network structure has an important influence on the success of the search process and in the average utility of the system. Scale-Free structural parameters facilitate the success of the search process because their structural properties increase the number of agents that can be reached. The experiments also show that even though there are certain values of the reward that are enough to promote cooperation, these values are not enough to obtain a positive average utility value.

Acknowledgements

This work is supported by TIN2011-27652-C03-01 and TIN2012-36586-C03-01 projects and PROMETEOII/2013/019.

[1] L. Backstrom and J. Leskovec. Supervised random walks: Predicting and recommending links in social networks. 2011.

[2] X. Ban, J. Gao, and A. van de Rijt. Navigation in real-world complex networks through embedding in latent spaces. In *Meeting on Algorithm Engineering and Experiments (ALENEX10)*, pages 138–148, 2010.

[3] D. O. Cajueiro. Optimal navigation in complex networks. *Physical Review E*, 79(4), 2009.

[4] Y. Chawathe, S. Ratnasamy, L. Breslau, N. Lanham, and S. Shenker. Making gnutella-like p2p systems scalable. In *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, SIGCOMM '03, pages 407–418. ACM, 2003.

[5] L. F. Costa, O. N. Oliveira Jr., G. Travieso, F. A. Rodrigues, P. R. Villas Boas, L. Antiqueira, M. P. Viana, and L. E. C. Rocha. Analyzing and modeling real-world phenomena with complex networks: a survey of applications. *Advances in Physics*, 60(3):329–412, 2011.

[6] E. Del Val, M. Rebollo, and V. Botti. Promoting Cooperation in Service-Oriented MAS through Social Plasticity and Incentives. *Journal of Systems and Software*, 86(2):520–537, 2013.

[7] E. Del Val, M. Rebollo, and V. Botti. Enhancing Decentralized Service Discovery in Open Service-Oriented Multi-Agent Systems. *Journal of Autonomous Agents and Multi-Agent Systems*, 28(1):1–30, 2014.

[8] A. Fronczak and P. Fronczak. Biased random walks in complex networks: The role of local navigation rules. *Phys. Rev. E*, 80:016107, Jul 2009.

[9] D. Fudenberg and E. Maskin. The Folk Theorem in Repeated Games with Discounting or with Incomplete Information. *Econometrica*, 54(3):533–554, 1986.

[10] C. Gkantsidis, M. Mihail, and A. Saberi. Random walks in peer-to-peer networks: Algorithms and evaluation. *Performance Evaluation*, 63(3):241–263, March 2006.

[11] L.-M. Hofmann, N. Chakraborty, and K. Sycara. The evolution of cooperation in self-interested agent societies: A critical study. In *The*

10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2, AAMAS '11, pages 685–692, Richland, SC, 2011. International Foundation for Autonomous Agents and Multiagent Systems.

[12] H. Janzadeh, K. Fayazbakhsh, M. Dehghan, and M. S. Fallah. A secure credit-based cooperation stimulating mechanism for manets using hash chains. *Future Gener. Comput. Syst.*, 25(8):926–934, Sept. 2009.

[13] J. J. Jaramillo and R. Srikant. A game theory based reputation mechanism to incentivize cooperation in wireless ad hoc networks. *Ad Hoc Networks*, 8(4):416–429, 2010.

[14] J. Kleinberg. Complex networks and decentralized search algorithms. In *Proceedings of the International Congress of Mathematicians (ICM)*, 2006.

[15] S. Lee, S.-H. Yook, and Y. Kim. Searching method through biased random walks on complex networks. *Phys. Rev. E*, 80, 2009.

[16] Z. Li and H. Shen. Game-theoretic analysis of cooperation incentive strategies in mobile ad hoc networks. *IEEE Transactions on Mobile Computing*, 11(8):1287–1303, Aug. 2012.

[17] A. L. Lopes and L. M. Botelho. Improving multi-agent based resource coordination in peer-to-peer networks. *Journal of Networks*, 3:38–47, 2008.

[18] M. Luck, P. McBurney, O. Shehory, and S. Willmott. *Agent Technology: Computing as Interaction (A Roadmap for Agent Based Computing)*. AgentLink, 2005.

[19] Q. Lv, P. Cao, E. Cohen, K. Li, and S. Shenker. Search and replication in unstructured peer-to-peer networks. In *Proceedings of the 16th international conference on Supercomputing*, ICS '02, pages 84–95, New York, NY, USA, 2002. ACM.

[20] A. B. MacKenzie and L. A. DaSilva. *Game Theory for Wireless Engineers*. Synthesis Lectures on Communications. Morgan & Claypool Publishers, 2006.

[21] M. E. J. Newman. The structure and function of complex networks. *SIAM REVIEW*, 45:167–256, 2003.

[22] M. E. J. Newman. Resource letter CS-1: Complex systems. *American Journal of Physics*, 79(8):800–810, 2011.

[23] J. D. Noh and H. Rieger. Random walks on complex networks. *Physical Review Letters*, 92:118701, 2004.

[24] H. Ohtsuki, C. Hauert, E. Lieberman, and M. A. Nowak. A simple rule for the evolution of cooperation on graphs and social networks. *Nature*, 441(7092):502–505, 2006.

[25] C.-L. Pu and W.-J. Pei. Mixing search on complex networks. *Physica A: Statistical Mechanics and its Applications*, 389(3):587 – 594, 2010.

[26] J. M. Pujol, J. Delgado, R. Sangüesa, and A. Flache. The role of clustering on the emergence of efficient social conventions. In *IJCAI*, pages 965–970, 2005.

[27] A. P. Roberts and C. P. Haynes. Electrostatic approximation of source-to-target mean first-passage times on networks. *Phys. Rev. E*, 83:031113, Mar 2011.

[28] C. Sierra, V. Botti, and S. Ossowski. Agreement computing. *KI-Künstliche Intelligenz*, 25(1):57–61, 2011.

[29] V. Srivastava, J. Neel, A. B. MacKenzie, R. Menon, L. A. Dasilva, J. E. Hicks, J. H. Reed, and R. P. Gilles. Using game theory to analyze wireless ad hoc networks. *Communications Surveys & Tutorials, IEEE*, 7(4):46–56, 2005.

[30] V. Tejedor, O. Bénichou, and R. Voituriez. Close or connected: Distance and connectivity effects on transport in networks. *Physical Review E*, 83(6), 2011.

[31] B. Yang and H. Garcia-Molina. Efficient search in peer-to-peer networks. In *Proceedings of the International Conference on Distributed Computing Systems (ICDCS)*, 2002.

[32] Z. Zhang, A. Julaiti, B. Hou, H. Zhang, and G. Chen. Mean first-passage time for random walks on undirected networks. *The European Physical Journal B*, 84:691–697, 2011.

[33] M. Zhong. Popularity-biased random walks for peer-to-peer search under the square-root principle. In *Proceedings of the 5th International Workshop on Peer-to-Peer Systems (IPTPS)*, 2006.

[34] T. Zhou. Mixing navigation on networks. *Physica A: Statistical Mechanics and its Applications*, 387(12):3025 – 3032, 2008.