

A Self-organized Mechanism of Resource Allocation in P2P Systems

Golnaz Vakili¹, Siavash Khorsandi¹, Thanasis G. Papaioannou²

¹ *Department of Computer Engineering and IT
Amirkabir University of Technology
Tehran, Iran*

{g-vakili, khorsandi}@aut.ac.ir

² *School of Computer and Communication Sciences
Ecole Polytechnique Fédérale de Lausanne (EPFL)
Lausanne, Switzerland
thanasis.papaioannou@epfl.ch*

Abstract

Since the performance of a system of interacting peers depends strongly on their individual contributions, the ratio between their provided and consumed resources should be set in line with the social welfare improvement, without relying on a central coordinator. In this paper, we devise a self-organized mechanism for cooperation policy setting of the interacting peers based on decision-theoretic analysis. By extensive simulation experiments, we demonstrate that when the proposed mechanism is followed, a Pareto optimal equilibrium emerges in the system from fairly coordinated decisions of the rational peers on their cooperation policies.

Keywords: distributed decision making; Pareto efficiency; fairness; particle swarm optimization; rational peers;

1. Introduction

As opposed to traditional client-server networking, the peer-to-peer paradigm is flexible and distributed in nature and allows individual peers (or users) to play the roles of both consumer and provider at the same time. Due to these characteristics, it has been widely deployed for file sharing and more recently for multimedia streaming systems. The performance of existing peer-to-peer systems relies on the cooperation of the peers and the contribution of their resources^{1,2,3,4,5}. However, it can be highly variable and unpredictable as there is no central authority to coordinate resource contributions of individual peers.

To improve the overall performance of the system, the rational peers should be incentivized to set the ratio between their provided and consumed resources in a self-organized manner. Two classes of incentive schemes

have been considered in the literature⁶: soft and hard schemes. Hard schemes, which are monetary-based, includes various methods of token-exchange and micro-payment, while in soft schemes, fairness of cooperation is the main concern and it is promoted by prioritizing peers according to their cooperation level. Under either of these incentive schemes, many different mechanisms of resource allocation^{4,7,8} and policies of server selection^{9,10,11} have been proposed, to respectively lead the decisions of rational peers as the provider and the consumer of resources. While they successfully provide incentives by differentiating the peers according to their contributions, the adjustment of their sharing ratio (the ratio between their upload and download resources) in accordance with the inherent dynamics of peer-to-peer systems is not explicitly considered. In other words, it is as-

sumed in the aforementioned works that each peer in the system has separate capacities for upload and download (as in the case of ADSL), and hence the upload capacity is the only possible bottleneck. However, this assumption is not realized in many systems where the access technology does not separate upstream and downstream flows (such as wireless or Ethernet-like protocols)¹². Strictly speaking, each peer has to decide on how to share its limited-capacity resources between upload and download simultaneously.

We focus in particular on setting the cooperation policies of individual peers by adjusting the ratio between their upload and download resources under an incentive scheme. Otherwise, rational peers would prefer to exploit their full capacity as the download resource in order to maximize their individual utilities, and the system would collapse. To this end, we devise a self-organized interaction mechanism among the participating peers by taking a decision-analytic approach¹³. The motivation is that under an incentive scheme, the interaction mechanism can be formulated as a process by which individual peers, observe the other participants' behavior, learn from these observations, and make the best decision in response to what they have learned. More specifically, interacting participants of a peer-to-peer system, which adjust their cooperation policies concurrently and adaptively, exhibit general properties of an individual-based Lagrangian swarm model¹⁴, since:

- the system consists of many individual peers with similar and simple functionalities;
- the interactions among the peers are based on simple behavioral rules by exploiting only local information that is exchanged directly via the environment;
- the overall behavior of the system results from the interactions of the individual peers with each other, which is emergence;
- the interactions of peers are realized in a distributed manner without relying on a centralized coordinator, that is self-organization.

In the context of a peer-to-peer system, e.g. a file-sharing or a media-streaming system, the behavior of an individual peer is defined by its cooperation policy which is characterized by the ratio between its upload and download resources (e.g. storage space in file-sharing system or bandwidth in media-streaming system). Moreover, the

social welfare is defined as the overall behavior of the system that is desired to emerge from fairly coordinated cooperation policies of individual peers.

Therefore, in order to develop the interaction mechanism in a self-organized manner, we model the cooperation policy setting as a swarm-based decision making process where distributed peers are represented by individual particles in the swarm. By means of extensive simulation experiments, we demonstrate that by taking this approach, the proposed mechanism results in fairness and Pareto-optimality in the system, and while rational peers set their cooperation policies through an iterative learning process to maximize their own utilities, their decisions are coordinated in line with the social welfare of the system as well.

This paper is organized as follows: In Section II, we develop the model of the system by describing the expected behavior of its constituent peers. In Section III, we propose our decision-analytic approach toward cooperation policy setting and then we evaluate the proposed approach through a simulated experimental setup in Section IV. Finally, after exploring some related works in Section V, we conclude our work in Section VI.

2. System Model

We consider a peer-to-peer system, e.g. a file-sharing or a media-streaming system, that consists of a set of interacting peers who participate in the system by contributing their resources. The interaction of participating peers for cooperation policy setting can be defined as a *stochastic game* $G = \langle N, \{D_i\}_{i \in N}, X, P, \{U_i\}_{i \in N} \rangle$ where N is the set of peers, D_i is the finite discrete set of actions (or cooperation policies) available to peer i , X is the finite discrete set of system states that can be defined with respect to the cooperation policies of the constituent peers, P is the state transition probability of the system, and U_i is the reward function of peer i for contributing to the system.

Generally, in such multi-player game, the state transitions and consequently the rewards are the result of the joint actions of all the players. However, in a peer-to-peer system, the reward of a peer is strongly dependent on the cooperation policies of the peers that are included in its environment and the peer benefits due to their participation. In other words, each peer has its own finite state space X_i corresponding to its environment, and hence it is associated with a local state transition function; where

$P_{d_i}^{x_i x'_i}$ denotes the probability of transitioning from a state $x_i \in X_i$ to another state $x'_i \in X_i$ by taking action $d_i \in D_i$ at x_i . Through their interactions, the rational peers should find the optimal policies in order to maximize their rewards.

We consider the state space of the system, X , as the set of all possible outcomes of this stochastic game. Let U_i^x be the payoff to peer p_i in outcome $x \in X$. The social welfare of the system in state x is defined as the sum of individual payoffs $\sum_{i \in N} U_i^x$ under this outcome. In order to improve the social welfare of the system, Pareto optimality and fairness are defined as the two necessary requirements for the mechanism of cooperation policy setting. The system is in a Pareto optimal state x if and only if for any other outcome preferred by some peer, another peer will prefer the current outcome:

$$\forall x' \in X, \exists i \in N; U_i^x \geq U_i^{x'} \quad (1)$$

The requirement of fairness is also fulfilled if and only if the maximum difference of payoffs under the outcome x is minimized:

$$\forall x' \in X, \max_{\forall i, j \in N} \left| \frac{U_i^x}{b_i} - \frac{U_j^x}{b_j} \right| \leq \max_{\forall i, j \in N} \left| \frac{U_i^{x'}}{b_i} - \frac{U_j^{x'}}{b_j} \right| \quad (2)$$

where b_i is the contribution of peer i to the system. In order to develop our swarm-based stochastic search method for this iterative learning process (in Section 3), we first model the system by describing the expected behavior of its constituent peers. A question that might be posed is why a quantitative model can give a reasonable description of the peers' behavior in the system. The fundamental results of decision theory directly address this question, by showing that any decision maker who is rational should always behave so as to maximize the mathematical expected value of some utility function, with respect to some subjective probability distribution¹³. That is, a rational strategic peer's behavior should be describable by a utility function, which gives a quantitative characterization of its rewards, and a subjective probability distribution, which characterizes its beliefs about other peers' behavior in its environment. Formally, the behavior of a rational peer i is given by:

$$E^{P_i}[U_i^{x_i}] = \sum_{d_i \in D_i} \sum_{x'_i \in X_i} P_{d_i}^{x_i x'_i} \cdot U_i(x'_i, d_i), \quad (3)$$

where P_i is the subjective probability distribution of local state transitions for the environment of peer i , and

$E^{P_i}[U_i^{x_i}]$ is the expected utility of peer i over the distribution P_i at state x_i . In order to construct a dynamic model based on the interaction of the participating peers, we first develop their individual reward function as fully rational expected utility maximizers and deal with the subjective probability distribution of their local state transitions in Section 3.

A peer i 's utility function under a soft incentive scheme is determined by its strategy choices and its local states, which are dependent on the benefit and cost of contributing to the system, as will be discussed in more detail.

2.1. Measuring the Cost and Benefit

In our setting, the strategy s_i that a peer i chooses to set its cooperation policy, reflects its decision on the changes in its cooperation level. In other words, a peer i 's cooperation policy (or action) is denoted by $d_i \in [0, 1]$ and can be defined as a numerical assessment of the ratio between the peer's upload and download resources in the system. The resources can be considered as the storage space in file-sharing systems or the bandwidth shared by participating peers in media-streaming systems. The definition of d_i is, nevertheless, acceptable as long as $\Delta d_i = d'_i - d_i$ (which represents peer i 's strategy s_i) can be quantified and treated as a decision variable.

For contributing to the system with cooperation level of d_i , a peer i incurs a cost that can be either in terms of the storage cost or the transmission cost in the aforementioned kinds of peer-to-peer systems. In case of a linear cost function which is commonly adopted in the literature^{15,16,11,6} the total cost could be assumed to be given by $c_i d_i$. However, more generally, the incurred cost is a function of cooperation policy and would be represented by $c(d_i)$.

On the other hand, the cooperation of each peer potentially benefits the other participating peers in the system and particularly in its environment. A participating peer j can be of benefit to i as a result of serving the required files of i in a file-sharing system or by means of supplying the streaming sessions of i in a media-streaming system. This benefit is represented by a matrix B , where $b_{ii} \equiv 0$, and b_{ij} denotes how much the cooperation of j is worth to i ; it can be measured as the relative buffering progress of peer j (e.g. in⁹), its service level (e.g. in^{4,7}), or its quality of service (e.g. in¹⁰), to the other participating peers. Obviously, the global contribution of peer i to the system, denoted by b_i , can be defined in terms of its local benefits to the other participants (e.g. in¹⁷). In general, $b_{ij} \geq 0$, and the set of peers with

$b_{ij} > 0$ is called the neighborhood of i . In other words, the neighborhood of a peer i , which is formally defined as $N_i = \{j \in N \mid b_{ij} \neq 0\}$, can be viewed as its local environment, in the sense that $N_i \subset N$ represents the subset of all the participants in the peer-to-peer system that i may interact with them directly.

2.2. Modeling the Incentive

Two prominent forms of soft incentive schemes that have been used widely^{3,4,18,19,20,15,11,21,22} are peer-approved and service-quality. They both use rating methods to assign points to each peer in the system according to its cooperation level. Under a peer-approved scheme, each peer only serves requests from peers with equal or higher ratings. In contrast, under a service-quality scheme, each peer serves all requests, but serves peers with higher ratings first. The authors in¹¹ applied a mechanism design approach to show that the service-quality scheme is a more promising form of modeling an optimal incentive mechanism in comparison with the peer-approved scheme.

In this paper, we aim at analyzing the interactions of peers under a soft incentive scheme, and hence we are primarily interested in modeling the effect of the incentive mechanism rather than implementing it. To this end, any reasonable function that is monotonically increasing in the cooperation policy of a peer can be considered. This is because, the benefit a peer can obtain from the system is proportional to its cooperation level under the service-quality form of a soft incentive scheme. Thus, we can simply model it as $f(d_i) = d_i$ or use a more complex function of d_i , since the exact form of the function is not relevant to our analysis.

2.3. Defining the Utility Function

As we discussed, the total benefit that a peer i can draw from the system is $\sum_{j \in N} b_{ij} \cdot d_j$ which is the weighted sum of the cooperation level of the other participating peers, while the weights reflect how much their cooperation is worth to i (according to their contributions). It is clear that the developed function in this general form is applicable to the settings where the environment of i consists of a limited set of peers that are included in its neighborhood; since $b_{ij} = 0$ for all the other $j \notin N_i$.

As already explained above, in order to take into account the effect of the incentive scheme, the expected benefit of i should be proportional to its cooperation level; thus, $f(d_i)$ - the function that is defined to model the incentive - is considered as a coefficient of the total benefit

in the utility function of i . Consequently, in each environment state x_i - that is defined by the cooperation policies of the participating peers in N_i - a nonnegative real-valued utility function U_i can be developed for i corresponding to the total expected benefit and the cost of contributing to the system. Strictly speaking, the total utility U_i that i will derive by setting the cooperation policy d_i in the state x_i is:

$$U_i(x_i, d_i) = f(d_i) \cdot \sum_{j \in N} b_{ij} \cdot d_j - c(d_i) \quad (4)$$

In the next section, we propose our decision-analytic approach to the cooperation policy setting, and then we evaluate how much the resultant individual decisions of the participating peers are in line with the social welfare of the system, with reference to Eq.(1) and Eq.(2).

3. The Decision-Analytic Approach

In this section, we propose a decision-analytic approach to devise the interaction mechanism of peers for cooperation policy setting under a soft incentive scheme. The decision-analytic approach to player i 's decision problem is to first assess some subjective probability distribution to summarize player i 's beliefs about which strategies will be used by the other players and then to select a strategy for player i that maximizes its expected utility with respect to these beliefs²³. However, there might be a fundamental difficulty in developing the decision-analytic approach; to assess its subjective probability distribution over the other players' strategies, player i may realize that the optimal strategies of the other players cannot be determined until their subjective probability distributions over player i 's possible strategies have been assessed. Thus, player i cannot predict the other players' behavior until it understands what they expect it rationally to do, which is of course, the problem it started with²⁴.

To overcome this difficulty, we predicate our decision-analytic approach to cooperation policy setting on an iterative learning process based on swarm intelligence²⁵. However, in order to formulate the notion of rationality and derive the implication of this notion for the emergent behavior, some modifications are necessary in modeling this process. The traditional model assumes that all particles in the swarm work together cooperatively to achieve a common goal; meanwhile, a peer-to-peer system consists of participants which are strategic and rational. In other words, they wish to maximize their own utility and hence they choose their strategies to achieve this goal. Thus, we make two modifications in order to construct a dynamic model for the process of

cooperation policy setting:

- Instead of a single global objective for all particles, distributed local objectives (utility functions) are defined for individual peers.
- The interaction of participating peers is represented as a non-cooperative game - each particle as a decision-maker wants to maximize its own utility.

Therefore in our approach, each peer makes an estimation of objective probability distribution^b over the other players' strategies. Meanwhile, individual opinion of the peer based on its past experience is also considered in this iterative decision making process. In other words, observable strategies of other peers are monitored by each peer in a sequence of iterations; then, based on this empirical evidence and its prior experience, each peer can decide rationally on a strategy in every iteration. Through this chain of decisions that are made based on a method inspired by particle swarm optimization²⁶, each participating peer concludes its final cooperation policy with respect to the other peers' behavior. It should be mentioned that particle swarm optimization is a stochastic evolutionary algorithm in which every particle in the swarm has a position in a socio-cognitive space; it updates its position in order to optimize an objective function, according to both social and cognitive behaviors. These behaviors are involved in the whole optimization process with two control parameters and respectively take the role of local and private guide of a particle in search of optimal value of corresponding objective function. With respect to the objective function, the best position that a particle has reached so far is called its private guide and the best position visited by one of its neighbors is the particle's local guide.

Strictly speaking, any particular peer i in the system, interacts only with a limited set of all possible peers, the ones that are included in its neighborhood N_i ; each peer i 's optimal policy should maximize its expected utility U_i with respect to the objective probability distribution over the possible policies of the other peers in N_i . To achieve this goal, each peer i sets its final cooperation policy through an iterative decision making process: in every iteration, each peer i monitors the strategies of the other peers in its neighborhood N_i , evaluates their strategies and chooses its strategy in the next iteration, s_i^{next} , with respect to the evaluation result and to its own experience (what it has learned up to the current iteration). Specifically, s_i^{next} is given by:

$$s_i^{next} = s_i^{curr} + r_1 c_1 (d_p - d_i^{curr}) + r_2 c_2 (d_n - d_i^{curr}) \quad (5)$$

^bObjective probabilities are determined based on a recorded observation rather than a subjective estimate.

where r_1 and r_2 are two distinct random values in $[0,1]$, c_1 and c_2 are the mentioned control parameters, determined theoretically in particle swarm optimization to properly balance between exploration and exploitation. d_p is the best previous cooperation policy of the peer itself and d_n denotes the best cooperation policy of all other peers in its neighborhood N_i . It should be noted that the cooperation policies are evaluated by each peer i according to the local utility function U_i defined in Eq.(4). Then, the cooperation policy d_i of peer i is revised by:

$$d_i^{next} = d_i^{curr} + s_i^{next} \quad (6)$$

As we mentioned, d_i^{next} would be a feasible cooperation policy for peer i , if the revision of d_i^{curr} results in a value between 0 and 1; otherwise d_i^{next} would be set to:

$$d_i^{next} = \begin{cases} 1 & ; d_i^{curr} + s_i^{next} > 1 \\ 0 & ; d_i^{curr} + s_i^{next} < 0 \end{cases}$$

Theorem 1. *The proposed mechanism for the cooperation policy setting converges to a Pareto optimal state in the system.*

Proof. By following the proposed mechanism, the iterative decision making process of interacting peers will converge when $\forall i \in N : s_i \rightarrow 0$ in successive iterations. In this situation, coordinated cooperation policies of participating peers in the system are led into a Pareto optimal operating point, since the mentioned condition of convergence is held when no peer can improve its utility by deriving from its current cooperation policy. In other words, for any other outcome preferred by some peer, another peer will prefer the current outcome. Then with reference to Eq.(1), it can be claimed that the operating point of the system has converged to a Pareto optimal state. \square

4. Evaluation Results

4.1. Simulation Model

In this section, we evaluate our self-organized approach to cooperation policy setting by studying the behavior of the interacting peers, through a set of simulation experiments. More specifically, we simulate a system of 1000 interacting peers who provide and consume bandwidth simultaneously for content sharing, and they are connected to a backbone network through their access

links with different capacities^c. They send and receive different number of requests among each other through their random overlay interactions. We evaluate the proposed mechanism of resource allocation by investigating individual decisions of the participating peers on the allocation of their limited capacities among upload and download bandwidth.

In all the experiments, the following heterogeneous setting is used to setup the simulation model unless some changes are explicitly mentioned for some of the scenarios. We choose the initial values of d_i assigned to peers $i \in N$ from a Uniform distribution over $[0,1]$. In practice, a peer i interacts only with a small subset of other peers in the system, and hence b_{ij} is non-zero only for a few values of j , specifically for the j that is included in N_i . Thus, the neighborhood size of individual peers is chosen to be in average 2% of N , while the non-zero values of b_{ij} are randomly selected between $(0,1]$.

The following form has been chosen to define the function that models the incentive in our experiments:

$$f(d_i) = \frac{d_i^\alpha}{d_i^\alpha + \beta} ; \alpha, \beta > 0, \quad (7)$$

where the choice of the exponent α determines how “step-function-like” the incentive model is. For small values of α for example $\alpha = 1$, the function is smooth; but for larger values, for example $\alpha = 10$, the function has a steep step. Thus, for large values of α , $f(d_i)$ - that converges to ideal step-function - can be used to model a peer-approved incentive mechanism; since in this case, each peer almost serves only requests from the peers with contributions equal or above the step; while the threshold of step is determined by the value of β . It is clear then, that for small values of α , $f(d_i)$ can model a service-quality mechanism in which a peer serves all requests of the participating peers, but only proportionally to their cooperation levels (these two forms of incentive schemes were discussed in Section 2.2). Here, we set $\alpha = \beta = 1$. Regarding the cost function, the linear form of $c(d_i) = c_i \cdot d_i$ is adopted in most of the experiments.

We study the proposed mechanism of resource allocation in this simulated setup, to evaluate its effectiveness towards fairness and Pareto optimality in line with the social welfare, by coordinating the cooperation policies of the participating peers.

^cThe capacity of the backbone network is regarded sufficient to provide the functionality of the peer-to-peer system. This assumption is quite realistic as previous studies have shown that the access links are most likely the bottleneck and not the backbone.

^dIt is defined as the ratio between U_i and b_i for peer i .

4.2. Performance Metrics

Our focus in this paper is how to dynamically allocate the resources (in this case the bandwidth) of rational peers in a self-organized manner to provide fairness and Pareto efficiency in the system. Therefore, we introduce a system-wide performance metric to capture these requirements based on the utilities of individual peers:

$$P = 1 - \frac{\frac{1}{N} \sum_{\forall i, j \in N} |U_i - U_j|}{\sum_{\forall i \in N} U_i} \quad (8)$$

where $0 \leq P \leq 1$. This metric is defined by the ratio between the fairness and the social welfare that emerge in the system, and hence it can measure and compare the performance of different resource allocation mechanisms with respect to the aforementioned requirements.

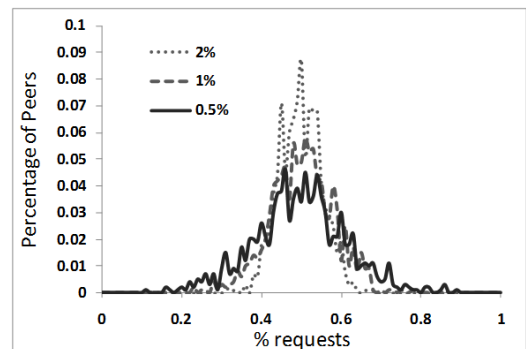


Fig. 1. Diversity of the percentage of requests that is received by individual peers.

Notice that the peers have various contributions to the system, as they send and receive different numbers of requests among each other through their random interactions (different probability distributions of the percentage of received requests by individual peers, that are simulated in our experimental setup, are illustrated in Fig.(1)). In order to take into account the variety of individual contributions in evaluating the performance, the normalized utilities of peers^d are used in Eq.(8).

4.3. Alternative Schemes

In order to evaluate the performance of our proposed mechanism, we compare it with other resource allocation schemes as follows.

4.3.1. Random Strategy

Under this scheme, peers follow a random strategy in determining the proportion of their capacity that is respectively assigned to provide and consume bandwidth. In other words, we consider that each peer i uses its capacity C_i to satisfy its own needs and offer the residual bandwidth to its requesters, while its demands vary randomly between $[0, C_i]$.

4.3.2. Evenly Averaged Policy

According to this policy, a peer i takes decision on how to allocate its upload and download resources based on the cooperation policies of its providers. To this end, it evenly averages over the resource allocation policies adopted by the peers $j \in N_i$ that are included in the neighborhood^e. According to the assessed cooperation policy, it determines the allocated bandwidth to its own requesters in turn.

To compare different resource allocation mechanisms, we denote our self-organized approach as RAM-S, while the random strategy and the evenly averaged policy schemes as RAM-R and RAM-E respectively when reporting the results.

4.4. Test Scenarios

In the first set of experiments, we consider a scenario in which a set of peers does not follow the resource allocation mechanism, to see how much the proposed mechanism is incentive-compatible and can promote the rational peers to be adherent. In each experiment, only a fraction of peers $\{0.1, 0.25, \dots, 0.85\}$ is considered to be participant in the proposed mechanism and the rest of them adopt the random strategy to allocate their resources.

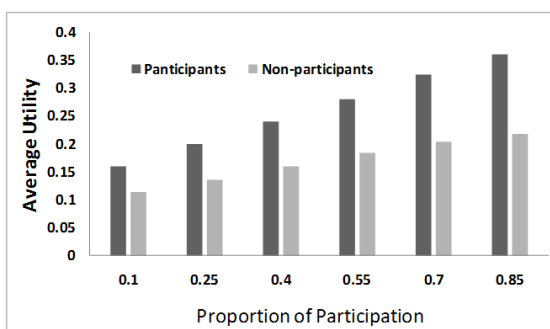


Fig. 2. Average utility of participating and non-participating peers.

^eThe cooperation policies of the peers can be averaged regardless of their different capacities, since they are in terms of their upload to download ratio.

As the illustrated results in Fig.(2) show, the participating peers achieve a higher utility in all the experiments, and the average utility improves as the proportion of participation increases in the system. The utility of individual peers versus the value of their contributions (in the last experiment where the proportion of participation is 0.85) is plotted in Fig.(3). This plot clearly shows that while there is a strong correlation between the utility of participating peers and their contributions, the achieved utility by the non-participants is not guaranteed at all. It means that a non-participating peer with a high value of contribution may gain a very low utility and vice versa.

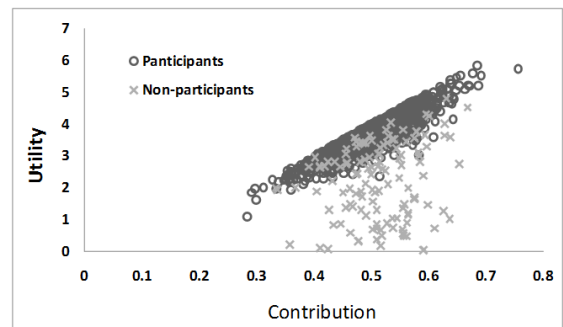


Fig. 3. Distribution of utility versus contribution for participants and non-participants.

In summary, this set of experiments demonstrated that, while the average utility of the non-participants is not as high as that of the peers that adhere to the mechanism, the value of their utilities is not proportional to their contribution levels as well. Thus, in the rest of the experiments, we can safely assume that all the peers in the system will follow the proposed mechanism of resource allocation.

In order to evaluate the performance of the proposed resource allocation mechanism (RAM-S), we should investigate how much it results in the fairness and improves the social welfare in the system. Meanwhile, we would like to see if the performance is acceptable regardless of the contribution value of the participating peers. Therefore, in the next set of experiments, we consider a scenario in which the contribution values of the peers are different in distribution, such that they have a certain mean value in each experiment; e.g. they have random values between $[0, 0.4]$ with the mean value of 0.2 in the first experiment and between $[0.6, 1]$ with the mean value of 0.8 in the last one. Then, the performance metric P , given in Eq.(8), is measured and compared for RAM-S, RAM-

E and RAM-R, as illustrated in Fig.(4). As expected, the social welfare and fairness will not improve by taking a random resource allocation strategy, and hence the performance of RAM-R is almost zero in all the experiments. Although RAM-S becomes more effective, while the contribution of peers to the system increases in average, it achieves a very high performance as well, even in the first experiment.




Fig. 4. Evaluation of the performance metric versus average value of individual contributions.

To scrutinize the results of this set of experiments, the distribution of the normalized utility of peers are also demonstrated in Fig.(5) for each experiment - the three markers show the maximum, average and minimum values of utilities resulting from each mechanism. It can be concluded that RAM-S improves the social welfare and the fairness more significantly in comparison with RAM-E and RAM-R, and it achieves a high performance even if the contributions of the peers are low in average.

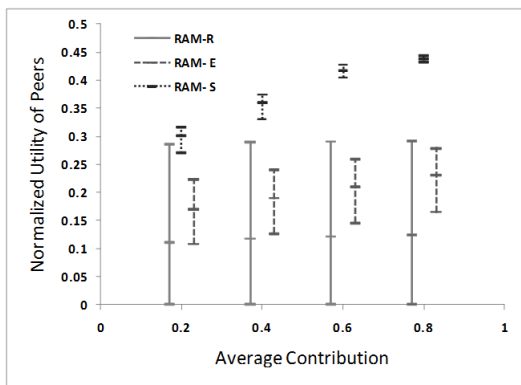


Fig. 5. Distribution of the normalized utility of individual peers.

As a complementary scenario to the previous one, we focus on the diversity of the contribution values of the peers, and increase it gradually in a set of experiments

to evaluate its effect on the performance. In the first experiment, we consider a homogeneous system of peers, in the sense that all the participants have the same contribution value (corresponding to the contribution similarity of one). In the next experiments, the mean contribution is fixed over this value (μ), while they are incrementally distributed more widely, such that in the last experiment the similarity of contributions is zero and they are distributed randomly between $[0, 2\mu]$. As illustrated in Fig.(6), while RAM-S outperforms the other schemes in all the experiments, it effectively improves the social welfare, and results in fairness, even if the similarity of contributions between heterogeneous peers is zero.

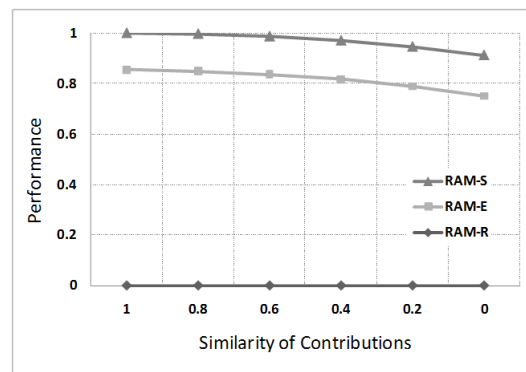


Fig. 6. Evaluation of the performance metric versus diversity of individual contributions.

In the next set of experiments, we investigate to what extent the proposed mechanism of resource allocation coordinates the cooperation policies of individual peers in a self-organized manner. Since the mechanism disseminates the state information of peers through their interactions, in the next scenario, our focus is on the number of connections formed by individual peers for their interactions and we change this parameter in each experiment.

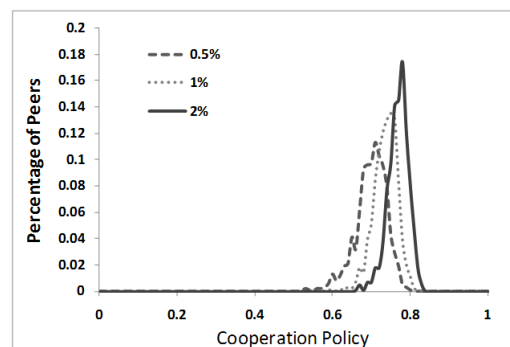


Fig. 7. Probability distribution of the cooperation policy, set by individual peers after 20 iterations.

To see the effect, the probability distribution of the cooperation policy set by the individual peers is demonstrated in Fig.(7), when they are connected to a fraction $\{0.005, 0.01, 0.02\}$ of the other participants. The PDF is assessed after 20 iterations in this set of experiments, and as observed, it is consistent for all of them. It only becomes denser with higher percentage of connections; thus, for smaller fractions, the approach takes longer to converge with no significant variation in behavior. It is important to note that the coordination is guaranteed by the proposed mechanism even if the individual peers interact with only 5 out of 1000 participants in the system.

To examine the self-organization of RAM-S more thoroughly, we measure the entropy of individual utility of the peers in the system. The resulting entropy from RAM-S, RAM-E and RAM-R is illustrated and compared in Fig.(8). The reduction in this mathematical measure is equivalent to self-organization, and as it can be observed, it is significantly reduced by RAM-S in all the experiments. Note that the entropy of the system is very close to 0 when individual peers interact with only 10% of the other participating peers.

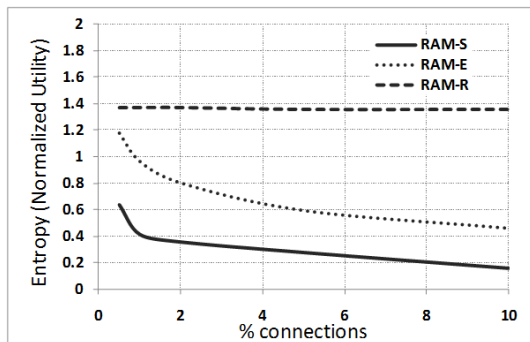


Fig. 8. The resulting entropy of individual utilities for different percentage of overlay connections.

The evolution of individual utilities versus the contribution values of the peers - in the experiment that each peer interacts with only 0.02 of the other peers - is plotted in Fig.(9). This figure shows that, by following the proposed interaction mechanism, the system of peers gradually converges to a completely fair state in which the achieved utility is very strongly correlated with the contribution (it almost scatters along the linear fitting line).

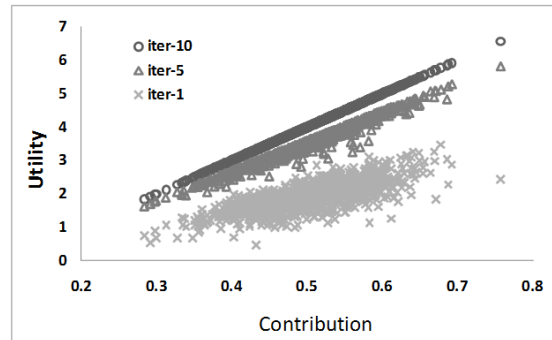


Fig. 9. The evolution of the utilities of individual peers versus their contribution values.

Finally, in the last set of experiments, we investigate the adaptability of our proposed approach to possible variation of the cost of resource sharing, e.g. due to dynamic availability of pure resource providers in the system. In other words, suppose that the cooperation policies of the peers have been currently set, but they have to be readjusted in order to adapt to the new cost of resource sharing in the system. In this scenario, we consider two random distributions between $[0, 0.2]$ and $[0.8, 1]$ as two extremes for initial cooperation policies to see whether in adaptation to a specific cost function, the results of the proposed mechanism in either of these cases are consistent (in terms of the coordinated cooperation policies of individual peers). The evolution of the cooperation policy distribution, in the both experiments, is demonstrated in Fig.(10) and Fig.(11) that are respectively corresponding to two different cost functions - the two markers show the maximum and minimum values of the distribution at each iteration.

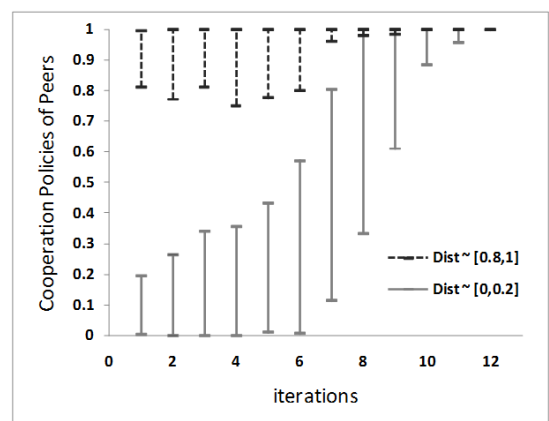


Fig. 10. The consistency of the results in adaption to a specific linear cost function.

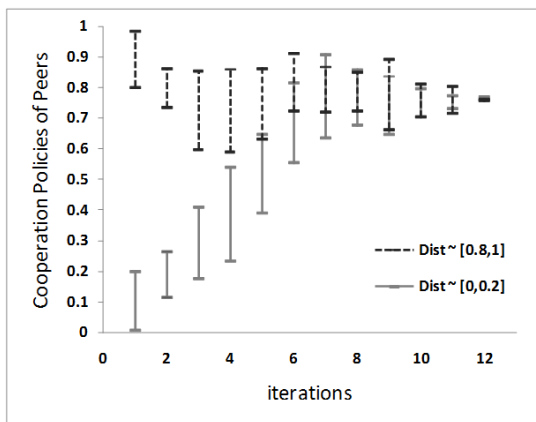


Fig. 11. The consistency of the results in adaption to a specific non-linear cost function.

According to this set of experiments, it can be concluded that, by following the proposed mechanism, the participating peers show an adaptive behavior and set their resource allocation policies in a consistent way regardless of their previously coordinated policies.

5. Related Work

Promoting fair cooperation of rational participants to improve social welfare in peer-to-peer systems has been extensively studied in the literature. Among these previous efforts we focus on a category that pays attention to cooperation policy setting. Specifically, they analyze a system of rational peers who decide on their contribution levels in order to maximize their individual utilities under an incentive scheme.

Two of the most recent approaches for cooperation policy setting are introduced in¹⁷ and²⁷. In the proposed mechanism in¹⁷, participating peers repeatedly interact to build up and maintain a dynamic social network based on the similarity of cooperation. Using the prisoner's dilemma problem to model net payoffs, they demonstrate that cooperation is incentivized by the mechanism in which it is positively correlated with payoff of the participants. This is because, in order to gain access to a social network of highly cooperative participants and a higher payoff (peers with higher levels of contribution gain higher payoff), a peer must itself exhibit at least that level of cooperation. This approach facilitates cooperation in decentralized systems where repeated bilateral transactions are required. However, peer-to-peer systems in general are characterized by large populations

and asymmetries of interest, resulting in few, if any, repeated transactions between any given pair of peers. The authors in²⁷ use a point-based mechanism to assign a global contribution (GC) value to the participants. The GC value characterizes the cooperation policy of a peer and reflects its resource (bandwidth) contribution to the entire network. The participating peers set their cooperation policy through transacting with each other. Based on the proposed rules for the calculation of the GC value, a peer may prefer transacting with a node with which the loss of GC is minimum or the gain of GC is maximum. In this approach, each peer needs to keep the history of its upload and download transactions with other peers and their GC values to calculate its GC value. Furthermore, to set a new cooperation policy through transacting with another node, a peer needs to know how they impact its GC. Thus, it should exchange a large amount of information (with reference to Eq.(13) in²⁷) with the other peer before making decision on doing transaction.

The authors in⁴ analyze the problem of designing optimal cooperation policies for individual peers in the presence of constraints, represented by an imposed upload-download fraction or sharing ratio. The constraints are enforced by a decentralized coalition of peers. However, to this end, they assume that the goal of the coalition is to maximize the efficiency of the system (the total amount of uploads and downloads in the system), which is not commonly assumed for rational participants of a peer-to-peer system. Then stochastic inventory management techniques are used to formulate cooperation policy setting as an optimization problem. The cooperation policy of individual peers is determined based on their desired service level in¹¹, through a peer selection mechanism. The proposed mechanism promotes higher contribution level of participants through differentiated quality of service using rank-order tournaments. Specifically, the cooperation policy of a participant is characterized as a score which in turn is mapped into a percentile rank, based on the global distribution of scores, so that a peer determines whether its rank among all participating peers in the system is sufficient to obtain an acceptable quality of service. Utility of a participating peer is a function of the quality of service it receives and the contribution cost. The quality of service depends on the rank ordering and hence the utility function takes into account the relative level of contributions of the participants.

As opposed to all the aforementioned works, we address the process of cooperation policy setting in peer-to-peer system where the participating peers have to decide on how to share their limited-capacity resources between

upload and download simultaneously. The only work, to the best of our knowledge, that has taken into account the adjustment of the ratio between upload and download resources of individual peers as a main concern is¹². The authors propose an incentive mechanism through which rational decisions of the peers are led to the operational point of system in which peers offer half of their resources for upload and they use the rest for download. However, the total upload and download in the system do not necessarily need to be in equilibrium for the system to operate efficiently⁴. Hence as opposed to¹², we explicitly define the ratio between upload and download resources of the peers as their cooperation policies to adjust them according to inherent dynamics of the system. Finally, our previous work²⁸ also dealt with the general problem of self-organized resource allocation, yet with a totally different approach. As opposed to²⁸ that totally depends on Reinforcement Learning to adjust peer cooperation policies based on reward values received from the environment, the developed approach in this present work is based on a stochastic direct search of the policy space, aiming to speed up the adaptation of the resource allocation to highly dynamic environments. While the approach in²⁸ is suitable for loosely coupled environments, our current approach is more applicable in tightly-coupled peer-to-peer systems.

6. Conclusion

To improve the performance of a peer-to-peer system which is highly dependent on the resource contributions from the participating peers, cooperation policies of the individual peers should be set and coordinated in a self-organized manner. Although many different mechanisms of resource allocation and policies of server selection have been proposed to coordinate rational decisions of the peers in this respect, they are specifically focused on the upload capacity as the only possible bottleneck. While, in many peer-to-peer systems, the adjustment of the sharing ratio of the peers (the ratio between their upload and download resources) in accordance with the inherent dynamics, can be a challenge that is not explicitly considered yet.

To this end, we devise a self-organized mechanism for cooperation policy setting of the interacting peers based on decision-theoretic analysis. More specifically, the cooperation policy setting is formulated as a process by which individual peers, observe the other participants behavior, learn from these observations, and make the best decision in response to what they have learned;

then, through their interactions, the rational decisions of the peers are led into the optimal policies (that maximize their utilities) based on a swarm-based optimization method. We demonstrate that the results of the proposed decision-analytic approach is in line with the social welfare of the system; since the constructed dynamic model converges to a Pareto optimal equilibrium, and this equilibrium emerges from fairly coordinated cooperation policies of the participating peers.

Acknowledgments

The first author would like to thank Prof. Karl Aberer for hosting her as a visiting researcher in the laboratory of Distributed Information Systems, Ecole Polytechnique Federale de Lausanne (EPFL) during part of the work involved in this paper.

References

1. M. Feldman and J. Chuang. Overcoming free-riding behavior in peer-to-peer systems. *ACM SIGecom Exchanges*, 5(4):41–50, 2005.
2. H. Park and M. van der Schaar. A framework for foresighted resource reciprocation in p2p networks. *IEEE Transactions on Multimedia*, 11(1):101–116, 2009.
3. X. Yang and G. Veciana. Performance of peer-to-peer networks - service capacity and role of resource sharing policies. *Performance Evaluation In P2P Computing Systems*, 63(3):175–194, 2006.
4. M. van der Schaar, D.S. Turaga, and R. Sood. Stochastic optimization for content sharing in p2p systems. *IEEE Transactions on Multimedia*, 10(1):132–144, 2008.
5. M. Sirivianos, X. Yang, and S. Jarecki. Robust and efficient incentives for cooperative content distribution. *IEEE/ACM Transactions on Networking*, 17(6):1766–1779, 2009.
6. O. Loginova, H. Lu, and X.H. Wang. Incentive schemes in peer-to-peer networks. *Theoretical Economics*, 9(1, Article 2), 2009.
7. H. Park and M. van der Schaar. Coalition-based resource reciprocation strategies for p2p multimedia broadcasting. *Transactions on Broadcasting*, 54(3):557–567, 2008.
8. R.T.B. Ma, S.C.M. Lee, C.S. Lui, and D.K.Y. Yau. Incentive and service differentiation in p2p networks: A game theoretic approach. *IEEE/ACM Transactions on Networking*, 14(5):978–991, 2006.
9. C. Liang, Z. Fu, Y. Liu, and C.W. Wu. Incentivized peer-assisted streaming for on-demand services. *IEEE Transactions on Parallel and Distributed Systems*, 21(9):1354–1367, 2010.
10. G. Tan and S.A. Jarvis. A payment-based incentive and service differentiation scheme for peer-to-peer streaming

- broadcast. *IEEE Transactions on Parallel and Distributed Systems*, 19(7):940–953, 2008.
11. A. Habib and J. Chuang. Service differentiated peer selection: An incentive mechanism for peer-to-peer media streaming. *IEEE Transactions on Multimedia*, 8(3):610–621, 2006.
 12. A. Satsiou and L. Tassiulas. Reputation-based resource allocation in p2p systems of rational users. *IEEE Transactions on Parallel and Distributed Systems*, 21(4):466–479, 2010.
 13. H. Raiffa. *Decision Analysis: Introductory Readings on Choices Under Uncertainty*. McGraw Hill, 1997.
 14. L. Edelstein-Keshet. Mathematical models of swarming and social aggregation. In *Proc. of International Symposium on Nonlinear Theory and its Applications*, pages 1–7, Japan, 2001.
 15. C. Buragohain, D. Agrawal, and S. Suri. A game theoretic framework for incentives in p2p systems. In *Proc. of International Conference on Peer-To-Peer Computing*, pages 48–56, Sweden, 2003.
 16. P. Antoniadis, C. Courcoubetis, and R. Mason. Comparing economic incentives in peer-to-peer networks. *Computer Networks*, 46(1):113–146, 2004.
 17. H. Nishida and Th. Nguyen. A global contribution approach to maintain fairness in p2p networks. *IEEE Transactions on Parallel and Distributed Systems*, 21(6):812–826, June 2010.
 18. K. Ranganathan. Incentive mechanisms for large collaborative resource sharing. In *Proc. of IEEE International Symposium on Cluster Computing and the Grid*, pages 1–8, Chicago, IL, USA, 2004.
 19. M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust incentive techniques for peer-to-peer networks. In *Proc. of ACM conference on Electronic commerce*, pages 102–111, NY, USA, 2004.
 20. K. Lai, M. Feldman, I. Stoica, and J. Chuang. Incentives for cooperation in peer-to-peer networks. In *Proc of the Workshop on Economics of Peer-to-Peer Systems*, CA, USA, 2003.
 21. P. Antoniadis, C. Courcoubetis, and B. Strulo. Incentives for content availability in memory-less peer-to-peer file sharing systems. *ACM SIG on Ecommerce*, 5(4):11–20, 2005.
 22. T.G. Papaioannou and G.D. Stamoulis. Reputation-based policies that provide the right incentives in peer-to-peer environments. *Computer Networks*, 50(4), 2006.
 23. H. Raiffa. *The Art and Science of Negotiation*. Harvard University Press, Cambridge, 1982.
 24. R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, 1991.
 25. J. Kennedy and R.C. Eberhart. *Swarm Intelligence*. Morgan Kaufmann Academic Press, 2001.
 26. J. Kennedy and R.C. Eberhart. Particle swarm optimization. In *Proc. of IEEE International Conference on Neural Networks*, pages 1942–1948, WA, Australia, 1995.
 27. S.M. Allen, G. Colombo, and R.M. Whitaker. Cooperation through self-similar social networks. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, 5(1, Article 4), February 2010.
 28. G. Vakili and S. Khorsandi. Coordination of cooperation policies in a peer-to-peer system using swarm-based rl. *Journal of Network and Computer Applications*, 35(2):713–722, 2012.