A Decision-Analytic Approach for P2P Cooperation Policy Setting

Golnaz Vakili[1], Thanasis G. Papaioannou[2], Siavash Khorsandi[1]
[1]Amirkabir University of Technology, Tehran, Iran
[2] Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

***Abstract*-While overall performance of peer-to-peer systems depends strongly on the amount of resource contributions made by individual peers, autonomous and rational peers make decisions on their cooperation policies (resource contributions) according to their individual utilities. To deal with the inherent conflict among individual utilities of the rational peers to improve overall performance of the system, we propose a decision-analytic approach that determines the appropriate cooperation policies of the individual peers in a distributed manner and coordinates their rational decisions in compliance with the social welfare improvement.**

Keywords- *rational peers; distributed decision making; Pareto efficiency; game theory;*

I. INTRODUCTION

The performance of a peer-to-peer system is highly dependent on the amount of resource contributions from individual participating peers [1-4]; however, it can be highly variable and unpredictable as there is no central authority to set resource contributions of peers or coordinate their cooperation policies. Thus, cooperation policies of individual peers (characterized by the amount of their resource contributions) should be set and coordinated in a distributed manner, such that the overall performance (i.e social welfare) of the system improves. However, autonomous and rational peers make decisions on their cooperation policies according to their own utilities, thus it is essential to incentivize participating peers to eliminate the inherent tension between individual utility of the peers and overall utility of the system.

Game theory has been used widely by previous works as a comprehensive analytical tool for study of incentive mechanisms in peer-to-peer systems. However, the traditional game-theoretic analysis lacks an explicit handling of the dynamics present in interactions among individual peers; e.g. it fails to represent the process by which one peer as a player, observes his opponent's behavior, learns from these observations, and makes the best move in response to what he has learned. Therefore, in this paper, we take an alternative approach based on decision-theoretic analysis [12] to model cooperation policy setting for participating peers of a peer-to-peer system. In our approach, each peer chooses its strategy according to observable strategies of the other peers. While through a swarm-based iterative learning process, rational peers set their cooperation policies such that maximize their own utility, their decisions are coordinated in a distributed manner to improve the social welfare of the system as well.

II. SYSTEM MODEL

Interacting participants of a peer-to-peer system that use simple rules to sequentially adjust their cooperation policies exhibit general properties of an individual based Lagrangian swarm model [9], as:

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

Therefore, cooperation policy setting for constituent peers of a peer-to-peer system is modeled as a swarm-based decision making process where distributed peers are represented by individual particles in the swarm. However, to adopt swarm model in the context of a system of autonomous peers, two modifications are necessary. The traditional model assumes that all particles in the system 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 made two modifications to adopt this model:

- Instead of a single common goal 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 wants to maximize its own utility.

Representing the interactions of peers as a non-cooperative game has been used widely in previous works such as [4-8]. As opposed to all of them that perform game-theoretic analysis to study the game, we propose an alternative decision-analytic approach to investigate the behavior of interacting players (peers) regarding the modified swarm-based model.

If it is assumed that N peers participate in the system, then the *utility* function of the ith peer will be defined as U_i. The behavior that a player adopts while interacting with other players is known as that player's *policy*. In our setting, a peer's policy is its level of cooperation and defined as d_i. A player chooses a *strategy*, defined as s_i, with respect to others to set its cooperation policy. Thus, the strategy of a peer reflects its decision on the change in its cooperation level (policy). A peer's utility is determined by its strategy choices and depends on several parameters which are discussed as follows.

A. Measuring the Cost and Benefit

A peer's p_i cooperation level, d_i, is defined as a numerical assessment of that peer's contributed resources to the system. Contributed resources can be the amount of disk space in storage overlays or other metrics such as the amount of bandwidth shared by participating peers in streaming overlays. The definition of d_i is acceptable as long as its fluctuation can be quantified and treated as a decision variable (which is p_i's strategy s_i).

For each unit of cooperation level, the peer incurs a cost c_i. So the total cost for participating in the system for a p_i with cooperation level of d_i will be $c_i d_i$ (linear cost is commonly used in the literature [8,15-17]). On the other hand, cooperation of each peer potentially benefits other participating peers in the system. This benefit is represented by a matrix B, where b_{ij} denotes how much the cooperation of p_j is worth to p_i; e.g. it can be measured as the inverse of latency between them. For instance, if p_i is not interested in p_j's cooperation because of the long physical distance between them, then $b_{ij} = 0$. In general, $b_{ij} \geq 0$, and $\forall i : b_{ii} = 0$.

B. Modeling the Incentive

Among different forms of incentives, the one which promotes fairness of cooperation by prioritizing peers according to their cooperation level (policy) has been used widely in previous works [3-8]. We are primarily interested in the effect of such an incentive form rather than implementing a possible mechanism; as the benefit a peer can draw from the system is proportional to its cooperation level in this form of incentive, any reasonable function that is a monotonically increasing function of the cooperation policy of a peer can be considered to model it.

C. Defining the Utility Function

At this point, why should we expect that such a simple quantitative model can give a reasonable description of peers' behavior? 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 distribution [12]. That is, any rational strategic peer's behavior should be describable by a utility function, which gives a quantitative characterization of its preferences for outcomes, and a subjective probability distribution, which characterizes its beliefs about other peers' behavior. According to the defined cost and benefit parameters, the total utility U_i that p_i will derive by setting cooperation policy d_i in the system is:

$$U_i = bc_i . \sum_{j \in N} b_{ij} . d_j - c_i d_i \quad ; b_{ii} \equiv 0 \tag{1}$$

where bc_i is the function that is defined to model individual incentives and can be considered as the benefit coefficient in utility function. U_i is a nonnegative real-valued function, which is developed corresponding to the total expected benefit (the first term) and the cost (the second term) of participants in the system. How can we deal with subjective probability distribution, will be discussed in the next section.

III. THE DECISION-ANALYTIC APPROACH

In traditional game-theoretic analysis that has been used widely in previous works, the usual approach is to analyze and solve the decision problems of all participating peers together, as a system of simultaneous equations in several unknowns. In contrast, the decision-analytic approach to player i's decision problem is to first access some subjective probability distribution to summarize i's beliefs about what strategies will be used by the other players and then to select a strategy for i that maximizes its expected utility with respect to these beliefs [11].

A. Declaration

The decision-analytic approach to i's decision problem is to try to predict the behavior of the others first and then to solve i's decision problem last. However, there might be a fundamental difficulty in implementing the decision-analytic approach; to access 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 i's possible strategies have been accessed. 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 with [10]. To overcome this difficulty we propose an iterative learning process modeled on swarm intelligence to implement the decision-analytic approach. In this process, objective probability distribution over the other players' strategies is taken by each peer as an estimation of the subjective probability distribution. Thus, observable strategies of other peers are monitored by each peer in a sequence of iterations. Based on this empirical evidence, 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 optimization [14], each participating peer concludes its final cooperation policy with respect to the other peers' behavior.

To give a rigorous declaration of our decision-analytic approach, let us consider the modified Lagrangian swarm model (discussed in section II) to investigate a system of N peers $\{p_i \mid i : 1, \ldots, N\}$. In this system, any particular peer p_i interacts only with a limited set of all possible peers, the ones that are included in its neighborhood N_i; each peer p_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. To achieve this goal, each peer p_i sets its final cooperation policy through an iterative decision making process: in every iteration, each peer p_i monitors the strategies of the other peers in its neighborhood N_i locally, evaluates their strategies and choose 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) as follows:

$$s_i^{next} = s_i^{current} + r_1 c_1 (d_p - d_i) + r_2 c_2 (d_n - d_i) \tag{2}$$

where r_1 and r_2 are two distinct random values in [0,1], c_1 and c_2 are the control parameters, 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 p_i according to the local utility function U_i defined for it. Then the cooperation policy d_i of peer p_i is revised as follows:

$$d_i^{next} = d_i^{current} + s_i^{next} \quad (3)$$

If and when local utilities, U_is, which are defined for individual peers in the modified swarm model (equation 1), converge to the same value through this iterative learning process, the resulting cooperation policies will constitute the final set of decisions that maximize rational peers' utility in line with the overall utility of the system. In the following section, the analysis of the approach and the optimality of the final policies will be discussed in more details.

B. Analysis

After proposing a feasible decision-analytic approach for players, we employ Nash equilibrium analysis to investigate whether the predicted strategies for the participating peers by the decision-analytic approach form a Nash equilibrium of the game. To this end, as demonstrated in [8] for a similar quantitative model of the system, a game-theoretic approach can be deployed as follows. The following form has been chosen to define the function that models incentive:

$$benefit\ coefficient: bc_i = \frac{d_i^{\alpha}}{1 + d_i^{\alpha}}; \quad \alpha > 0 \quad (4)$$

The choice of the exponent α determines how step-function-like the incentive function is.

First a simplified setting should be considered in which it is assumed that $b_{ij} = b; \forall i \neq j$; in other words in this system all peers derive equal benefit from everybody else. Therefore, by symmetry, the problem can be reduced to a 2-person game, which is analyzed by a similar methodology to the one used in Cournot duopoly model [13]. To find the set of fixed points that constitute Nash equilibrium, it is assumed that $\forall i: d_i = d$ that means all participating peers set the same cooperation policy simultaneously. In this case, the solution is:

$$\forall i: d_i = d^* = (b/2c - 1) \pm ((b/2c - 1)^2 - 1)^{1/2} \quad (5)$$

The results for the two peer system then applied to the N player system as well and therefore (for more details please refer to [8]):

$$d^* = (b(N-1)/2c - 1) \pm ((b(N-1)/2c - 1)^2 - 1)^{1/2} \quad (6)$$

In contrast with the game-theoretic approach, no limiting assumption is made in decision-analytic approach and hence the system is investigated in a more realistic setting. As we will numerically show, the above equilibrium is not formed when the cooperation policies are set based on the proposed decision-analytic approach. This is because the expected Nash equilibrium of the game is not the Pareto-optimal one, as the outcome derived from the decision-analytic approach based on modified swarm model would make all players better off.

To demonstrate that the tendency toward Pareto efficiency can be derived from our proposed approach, we perform numerical experiments in both homogeneous and heterogeneous settings. We choose the initial values of d_i from a standard Gaussian distribution. While in the game-theoretic approach it is assumed that all the participating peers set their cooperation level simultaneously to achieve Nash equilibrium, in the proposed approach the distribution d_i evolves every iteration and eventually converges to a Pareto efficient equilibrium. We choose the number of participating peers N to be from 500-1000. Since a peer p_i interacts only with a small subset of other peers, b_{ij} is non-zero only for a few values of j. The peers for which b_{ij} is non-zero are randomly selected. The size of the set for which $b_{ij} \neq 0$ is chosen to be 2 percent of N. In general, for smaller value of this fraction, the approach takes longer to converge.

We first consider a homogeneous setting in order to compare the experimental results of the proposed approach with the analytic results of the game-theoretic approach. If it is assumed that the system is completely homogeneous, the distribution of b_{ij} consists of a single peak at $b = b_{ave} / (N-1)$. As an instance, for b_{ave}= 5, based on the game-theoretic approach, the corresponding value of d^* from equation 6 would be 2.62; and thus the utility function value of peers would be 6.86, in Nash equilibrium (from equation 1). While in the same circumstance, the decision-analytic approach makes all peers better off and achieve a much higher U^* with the value of 15.61 (higher more than 50%).

We next consider a heterogeneous setting to investigate whether or not the results verify the results of homogeneous setting. In this setting, we choose the values of b_{ij} from a Gamma distribution such that b_{ave}= 5. The result of experiment shows that the values U_i^*s converge to the same value of 15.6 with a high confidence level. The value of α for all of the results is 1.0.

It is important to note that, in both homogeneous and heterogeneous settings, the mean of cooperation level of peers is much higher than their cooperation level in the game-theoretic approach (d_{ave}^* is around 40% higher than the value of d^* from equation 6). It demonstrates that the proposed approach improves the overall performance of the system in comparison with the game-theoretic approach. Maybe it seems that the final cooperation policy d_i that p_i sets, is not globally optimal as it is determined based on monitoring locally a limited set of peers N_i. But in fact, after p_i has set its cooperation policy, its strategy will be monitored by the peers it interacts with and those peers will adjust their own cooperation policy respectively. Therefore, the actions of any peer p_i will eventually reach all possible peers and the reactions of other peers to p_i's cooperation policy will affect p_i itself in the same way. Strictly speaking, through the proposed approach, globally optimal cooperation policies of peers emerge as a result of their local interactions.

In Figure 1, we show the average cooperation level of the peers as a function of average benefit. The lower trace, indicated by "□", is the solution from the analytic results of the game-theoretic approach. As expected, the cooperation level increases monotonically with increasing benefit while in comparison, it is improved significantly by the decision-analytic approach. Note that in this approach, the two sets of results for homogeneous and heterogeneous settings almost coincide with each other, which was expected for the average performance. In the *x* axis, the average benefit is scaled as $(b_{ave}/b_c)-1$ in which b_c is the critical value of benefit below which it is not profitable for a peer to join the system by the game-theoretic approach.

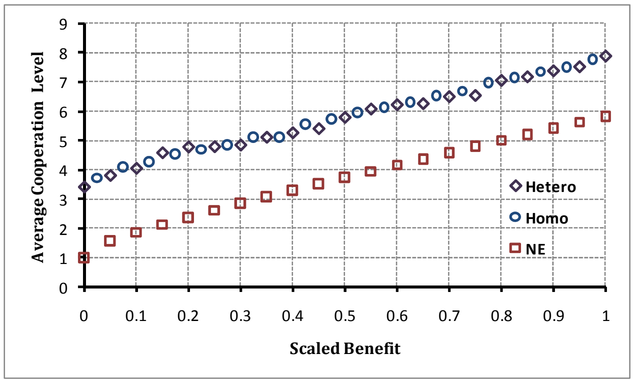

Figure 1. The comparison of the average cooperation level

In Figure 2 we show the convergence of the decision-analytic approach. The three data sets correspond to different values of average b_{ave}. As depicted, the convergence is fast and independent of the average value of b_{ave}. Similar to the game-theoretic approach, below a critical value of b_{ave} the iterations converge to zero that means the system collapses. It is interesting to note that the critical value of benefit obtained experimentally in this approach is less than 5% away from the calculated value of b_c by the game-theoretic approach. We have observed that for a wide set of initial conditions for d_i, after a sequence of iterations, the process tend to create a system in which most individual participants set the Pareto efficient cooperation policy.

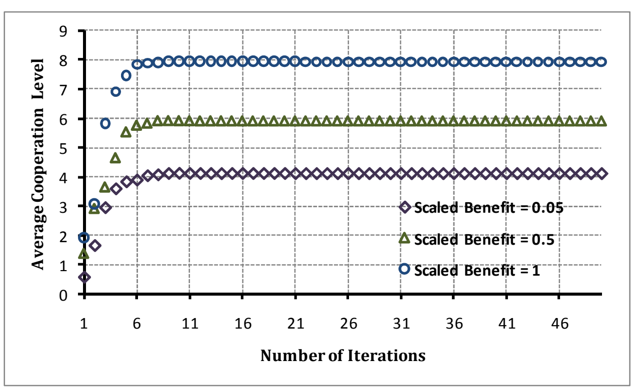

Figure 2. Average contribution level against number of iterations required to reach a set of Pareto efficient strategy

To conclude our work, it is demonstrated that by taking a decision-analytic approach based on the modified swarm model, rational decisions of the individual participating peers on their cooperation policies are adequately coordinated in a distributed manner and improve the overall performance of the system in both homogeneous and heterogeneous settings. Our approach quickly approximates a Pareto-optimal operating point of the system.

References

[1] M. Feldman, J. Chuang, "Overcoming Free-Riding Behavior in Peer-to-Peer Systems," in ACM SIGecom Exchanges, vol. 5, Issue 4, 2005, pp. 41 - 50 .

[2] H. Park, M. van der Schaar, "A Framework for Foresighted Resource Reciprocation in P2P Networks," in IEEE TRANSACTIONS ON MULTIMEDIA, vol. 11, no. 1, 2009, pp. 101-116.

[3] X. Yang, G. Veciana "Performance of peer-to-peer networks - service capacity and role of resource sharing policies," in Performance Evaluation In P2P Computing Systems, vol. 63, no. 3. (March 2006), pp. 175-194.

[4] M. van der Schaar, D.S.Turaga, R. Sood, "Stochastic Optimization for Content Sharing in P2P Systems," in IEEE TRANSACTIONS ON MULTIMEDIA, vol. 10, no. 1, 2008, pp. 132-144.

[5] K Ranganathan, "Incentive Mechanisms for Large Collaborative Resource Sharing," in Proceedings of IEEE International Symposium on Cluster Computing and the Grid,2004, pp. 1-8.

[6] M Feldman, K Lai, I Stoica, J Chuang "Robust Incentive Techniques for Peer-to-Peer Networks 04," in proceedings of ACM conference on Electronic commerce, USA, 2004, pp. 102-111.

[7] K Lai, M Feldman, I Stoica, J Chuang, "Incentives for cooperationin Peer-to-Peer Networks," in proceedings of the Workshop on Economics of Peer-to-Peer Systems, 2003, CA.

[8] C. Buragohain, D. Agrawal, and S. Suri, "A game theoretic framework for incentives in P2P systems," in proc. of Int. Conf. Peer-To-Peer Computing, Sweden, 2003, pp. 48-56.

[9] L. Edelstein-Keshet, "Mathematical models of swarming and social aggregation," in proceedings of International Symposium on Nonlinear Theory and its Applications, Japan, 2001, pp. 1-7.

[10] R. B. Myerson, *Game Theory: Analysis of Conflict*, Harvard University Press, Cambridge, 1991.

[11] H. Raiffa, *The Art and Science of Negotiation*, Harvard University Press, Cambridge, 1982.

[12] H. Raiffa, *Decision Analysis: Introductory Readings on Choices Under Uncertainty*, McGraw Hill, 1997.

[13] D. Fudenberg, J Tirole, *Game Theory*. MIT press, Cambridge MA, 1991.

[14] J. Kennedy, R.C. Eberhart, *Swarm Intelligence*, Morgan Kaufmann Academic Press, 2001.

[15] P. Antoniadis, C. Courcoubetis, and R. Mason, "Comparing Economic Incentives in Peer-to-Peer Networks," in Computer Networks, vol. 46, iss. 1, 2004, pp. 113- 146.

[16] A. Habib, and J. Chuang, "Service Differentiated Peer Selection: An Incentive Mechanism for Peer-to-Peer Media Streaming," in IEEE Transactions on Multimedia, vol. 8, no. 3, 2006, pp. 610–621.

[17] O. Loginova, H. Lu, and X.H. Wang, "Incentive Schemes in Peer-to-Peer Networks," in Theoretical Economics, vol. 9, iss. 1, 2009, Article 2.