On the Impact of Choice in Multi-Service P2P Grids

Álvaro Coêlho, Paulo Ditarso, Flavio Figueiredo, David Maia, Francisco Brasileiro

Universidade Federal de Campina Grande
Departamento de Sistemas e Computação
Laboratório de Sistemas Distribuídos
Av. Aprígio Veloso, s/n, Bloco CO
58.109-970, Campina Grande - PB, Brazil
Emails: {degas, pmaciel, flaviov, davidcmm, fubica}@lsd.ufcg.edu.br

***Abstract*—In this paper we consider a peer-to-peer grid system which provides multiple services to its users. In this system, an incentive mechanism promotes collaboration among peers. It has been shown that the use of a reciprocation mechanism in such a system is able to prevent free riding and, at the same time, promotes the clustering of peers that have mutually profitable interactions. However, when peers are subject to a budget limitation, each peer must select a subset of all services that can possibly be offered. In this work we show that the received utility is strongly dependent of the offered services. The main contributions of this work are a methodology to evaluate the impact of service changes in the obtained utility and how much different sets of offered services impact in the peer's utility. These results indicate that further research is needed, particularly for the development of heuristics to choose the best services to offer.**

I. Introduction

With the popularity of the Internet, peer-to-peer (P2P) systems are becoming an interesting way to obtain files, storage, processing and many other computational resources with high levels of availability.

P2P systems can be defined as networks of computers in which participants play both the role of a client and that of a server [1]. This is to say that wherever a new client is added, there is a new server that is also added to the system. In these systems, peers who are acting as servers must schedule the use of their resources in order to provide services to the other peers that are acting as clients at that time.

This scheduling process has been often done using two approaches: market-based and sharing-based [2]. Market-based approaches have been used in many systems [3], [4], [5]. Although market-based solutions give more strict guarantees on the quality of service provided, the nature of this approach implies in higher transaction costs for participating in such economy. This occurs because this kind of solution relies on the existence of a currency distribution system, banking services, auditing, and accurate pricing [6]. On the other hand, no currency schemes or any trustworthy central institution are used in sharing systems, and peers exchange resources based on a reciprocation scheme. In this environment, peers can make use of social mechanisms for monitoring and enforcement, because the information is loosely structured (and therefore, easier to obtain). So, this kind of solution decreases the transaction costs [7]. Reciprocation between peers is achieved through the use of incentive mechanisms.

Reciprocation-based mechanisms have been explored in previous works [6], [8], [9]. These mechanisms presume a single good to be exchanged, making the reciprocation schema simpler. However, they are not appropriate for the case in which peers need to exchange multiple goods (or services). In this case some additional considerations must be made, because peers may exchange a kind of service by another and peers may value services differently [10]. Consequently, the partnership between peers is based not only in the peer's behavior but also in the profitability of their interactions.

The Extended Network of Favors (ExtNoF) mechanism was proposed in [10] as a variation of the original Network of Favors. ExtNoF allows peers to exchange multiple services using the past interactions to find the more profitable partnerships, generating clusters of peers having compatible affinities that lead to mutually profitable interactions.

Services can be thought at the resource level (e.g. CPU, disk, bandwidth, etc), but also at a higher level, such as the execution of a particular software. Given the multitude of services that can be possibly offered, and the fact that peers have a limited budget, not all services can be offered by all peers. This imposes the need for each peer to choose a set of services to offer among all possible services. Each possible set has a specific cost and returns different profits.

In this paper we make an analysis of a P2P system in which multiple services are exchanged. We show that, for each peer, there are different set of services to be offered and each set gives different results. Moreover, we show also the impact of different choices of services in the peers' final utilities.

The rest of the paper is organized as follows. Related work to ours is discussed in Section II. Section III gives a formal description of the problem, and describes a utility function for the multi-service P2P grid system. A numerical evaluation of the system is presented in Section IV, in which different scenarios were considered. Finally, Section V closes the paper with our concluding remarks and perspectives of future work.

978-1-4244-2191-6/08/$25.00 ©2008 IEEE

II. Related Work

Systems where peers can exchange multiple services have been recently studied in different works. In one way, architectures are being proposed in order to support the exchange of services. Other works propose mechanisms and frameworks to incentive the reciprocation in P2P environments with multiple services.

In [11] the authors describe an architecture in which pools of machines can be donated by peers using a cluster manager called Cluster-On-Demand (COD). This architecture enables peers to donate their idle machines, defining how many machines must be donated to each other. The set of machines received by peers is perceived as a "virtual cluster", which is a pool of configured machines that are donated by different peers, allowing the receiver to use it as a real cluster. Moreover, the donating peer is able to configure the machines with multiple services, according to the receiver needs. The scheduling policy must be defined previously by the peers, and the services are configured in the machines after that. However, this work does not take into account the variation on the "utility" delivered by the grid when different set of services are offered by the peers. Furthermore, our work takes a business-driven management approach by considering a budget limitation for offering services.

In [12] a framework is proposed in order to manage exchanges in P2P systems. In this framework a peer uses a Reputation Vector for each peer with information about each service that it provides, in order to know how trustworthy the others are. The reputation of the peers is related to specific services, being defined by the ratio between the amount of service provided by the peer and the total amount of requested service. Using these reputation vectors, peers can infer how trustworthy are the other ones, and decide when and how much resources should be donated to them. They also define the overall local reputation vector as the sum of all services reputation for each peer. This work allows two policies of service allocation: the Trust-Based and the Enhanced Trust-Based Policy. While the first one is applied in environments where peers generate requests with the same rate, the second one is more useful when some peers produce many more requests than others. The proposed framework fosters trust between peers exchanging multiple services by consider some features as capacity and demand. However, there is no mechanism to help peers to find out how much the selection of services impacts on their profits, under a scenario with different values of services and a budget limitation.

The Network of Favors (NoF) is a reciprocation scheme that does not rely on the existence of banking services, trust or previous negotiation among peers. This mechanism is similar to the *tit-for-tat* scheme used in BitTorrent systems in which peers exchange pieces of a file based in their past bilateral interactions [9]. Additionally, the authors have proposed an extended version of the Network of Favors (Extended NoF - ExtNoF) in order to deal with multiple services [10]. ExtNoF is a reciprocation-based mechanism to encourage donation in peer-to-peer grids, and, unlike the single service version of the NoF, multiple services are shared explicitly, such as processing power, data transfer and storage. Both NoF and ExtNoF mechanisms address the problem of discourage free-riding and enable peers to find others who are mutually profitable to exchange services, even considering that peers value services differently. However, this work assumes that peers can often offer all the services, which is not realistic in some cases where the provision of services is submitted to financial (budget) and hardware limitations.

III. Problem Statement and System Model

In this section we present a formal description of the problem and the utility function of the scenario sketched in the Introduction.

A. Problem Statement

In a peer-to-peer grid, each peer can provide a set of different services such as processing power, storage, database management, etc. In the same manner, each service has a cost that depends on many factors as, for example, resource costs, kind of services, infrastructure installation and maintenance, etc. So, services probably will have different costs to different peers. All of this imposes a limitation to each peer, given that the cost to provide the whole set of services could be prohibitive. In this case, it is necessary to choose a subset of services to provide, that is limited by a budget. Additionally, since a peer must have an incentive to donate services to the others, a reciprocation scheme that guarantees this cooperation is essential.

Once peers cannot provide all the services due to the budget limitation, the selection of services to be offered is an interesting research problem to be investigated. It is necessary to explore what is the impact of multiple choices in the "utility" offered by the grid. It is also important to note that when a peer selects a set of service the cost is immediate, but the computing units will be received in the future. So, the problem can be defined as to find which set of services maximizes the total profit obtained. However, before that, it is essential to know what is the influence of different choices of services in the "utility" provided by the grid.

In this paper, we evaluate how the service selection impacts in the costs of a peer and in the expected amount of computing units reclaimed from the grid, according to the ExtNoF mechanism.

B. System Model

We consider that the time is slotted in T turns, and the duration of each turn t ($\in T$) is enough to make management decisions. We assume N peers sharing spare resources in a P2P grid, and a set of n services that could be offered by the peers. Each peer chooses which services to offer to the others according to a limited budget. So, we consider that peer i has a limited budget $\mathcal{B}_i$ to set and maintain the services. Moreover, the tuple $\mathcal{S}_i^t = \langle s_{i,1}^t, s_{i,2}^t, ..., s_{i,n}^t \rangle$ represents the set of services offered by peer i in a turn

t, with $s^t_{i,j} \in \{0,1\} \quad \forall \quad i \in \{1,...,N\}, \quad j \in \{1,...,n\}$ and $t \in \{1,...,T\}$. Assuming that peers may value services differently, the tuple $\mathcal{C}_i = \langle c_{i,1}, c_{i,2}, ..., c_{i,n} \rangle$ represents the peer i's cost per turn to offer each service. We also assume that these costs are constant over time.

In this P2P grid, units of computing power configured with specific services are exchanged by the peers. So, q_i represents the peer i's capacity in number of computing units. For the sake of simplicity, we consider that each peer i has a constant hardware cost per computing unit in each turn, represented by c^{hw}_i. Given all that, the peer i's total cost per turn is given by:

$$C^t_i = q_i \cdot c^{hw}_i + \sum_{j=1}^{n} s^t_{i,j} \cdot c_{i,j} \leq \mathcal{B}^r_i,$$

where $\mathcal{B}^r_i = \mathcal{B}_i - \sum_{k=0}^{t} C^k_i$ is the remaining budget till turn t.

We consider that each peer has a typical demand which represents the average amount of computing units and the average proportion of services typically required. Hereafter, we assume that the average amount of computing units requested by peer i is represented by d_i, and the average proportion of services is its *typical favor*, represented by the tuple $F_i = \langle f_{i,1}, f_{i,2}, ..., f_{i,n} \rangle$, where $0 \leq f_{i,j} \leq 1 \ \ \forall \ \ i \in \{1,...,N\}$, and $j \in \{1,...,n\}$, and $\sum_{j=1}^{n} f_{i,j} = 1$.

It is reasonable to assume that the profit obtained by peer i for using service j is higher than the costs incurred by maintaining the service j. So, we consider a profit factor $\mu_{i,j}$, which represents how profitable is the usage of the service j by the peer i. Moreover, in order to compute the profit obtained by peer i for executing its typical demand, it is necessary to know the amount of resources reclaimed from the grid. The total amount of computing units that peer i receives from the grid for each service in a certain turn is represented by the tuple $\mathcal{R}^t_i = \langle r^t_{i,1}, r^t_{i,2}, ..., r^t_{i,n} \rangle$. Finally, the total utility received by peer i is given by:

$$U_i = \sum_{t=1}^{T} (P^t_i - C^t_i),$$

where the profit obtained per turn (P^t_i) is 0 if peer i is donating services to other peers, and otherwise is given by the total profit achieved by the peer, which takes into account the cost of whole infrastructure used by it (the local capacity plus the resources received from the grid) for executing its typical demand:

$$P^t_i = \sum_{j=1}^{n} \mu_{i,j} \cdot C^t_{i,j}(q_i \cdot f_{i,j} + r^t_{i,j}),$$

where $C^t_{i,j}(x) = x \cdot c^{hw}_i + s^t_{i,j} \cdot c_{i,j}$.

In our setting, the balance of past interactions among the peers is calculated according to the extended NoF [2]. We define $Bal^t_i(k)$ as the peer i's balance in relation to peer k, in the turn t, given by:

$$Bal^t_i(k) = \begin{cases} max\{0; Bal^t_i(k) - \sum_{j=1}^{n} r^t_{k,j}(i) \cdot (c_{i,j} + c^{hw}_i)\}, \\ \text{if } i \text{ is donating to } k\text{; and} \\ Bal^t_i(k) + \sum_{j=1}^{n} r^t_{i,j}(k) \cdot (c_{i,j} + c^{hw}_i), \\ \text{if } i \text{ is receiving from } k\text{;} \end{cases}$$

where $r^t_{k,j}(i)$ is the amount of computing units that peer k receives from peer i for the service j, in the turn t.

IV. Numerical Evaluation

We have developed a simulator to evaluate the model described above. In this section we present the simulation scenario evaluated, and the obtained results.

We consider the time as being slotted in turns, and at the end of a fixed number of turns the simulation ends. At each turn, the peers being simulated interact and exchange resources. During the simulations, peers can be in two distinct states: donating or consuming. We have considered a system with average contention, i.e. each peer will have a 50% probability of being donating or consuming. Whenever peers interact, they update their balances of the peer it interacted with, and also their utility. Moreover, the simulations were carried out with $N = 100$ peers and $n = 6$ services. For the sake of simplicity, we assume a P2P grid with homogeneous peers whose local in-house capacity are one computing unit ($q_i = 1$). Moreover, we assume a fixed budget per turn of $\mathcal{B}_i/T = 1$, and a peers' hardware cost of 10% of their budget ($c^{hw}_i = 0.1$). Service costs were randomly chosen between 0 and 0.5, i.e. $0 \leq c_{i,j} \leq 0.5 \quad \forall \quad i \in \{1,...,N\}$ and $j \in \{1,...,n\}$. In this setting, 10% of a peer's budget is spent with hardware costs and the software cost can reach 90% of its budget ($\sum_{j=1}^{n} s^t_{i,j} \cdot c_{i,j} \leq 0.9$).

It is reasonable assume that the peers always offer the services that comprising their typical favor to the others, since these services are used to compute their local profit. We consider that the peers' typical favors are composed of two services. Given the fact that there are a total of six services in the system, peers can vary different possibilities of services to donate with the remaining four services, according to its budget. In order to understand the impact of these possible choices, we have varied the peer's demand (d_i) and profit factor ($\mu_{i,j}$). These two parameters have a direct impact on the utility obtained by each peer, since they increase the profit that a peer obtains from the grid. The values chosen for the demand were 2, 11, 51 and 101; that correspond to a peer requesting from the grid one, ten, fifty and one-hundred times its local capacity. When a peer requests 100 times its local capacity from the grid, we can say that it is eager because it can reclaim at most 99 computing units from the other peers. The profit factor was varied with values of 1.1, 2, 10, in order to cover different possibilities of results. So, we have carried out simulations for twelve different scenarios, varying all of these parameters.

We have evaluated the impact of choice on the final utility of peers considering the above configuration. In order to do such, we have varied all the possibilities of services to donate for each peer, and measured its utility considering this possibility. With a low profit factor ($\mu_{i,j} = 1.1$) every peer obtained only negative utilities. In the same manner, with a low demand $d_i = 2$ and a profit $\mu_{i,j} = 2$, we have found that some peers also reached only negative utilities. For the sake of simplicity for analyzing the results, we have unconsidered the scenarios in which peers obtained negative utilities. We can notice that, in these settings, peers with negative utilities would not remain in the system. So, the results explained in the following consider the other 7 scenarios in which peers have achieved positive utilities.

In the Figure 1 we demonstrate the *percentage increase* of utility considering the lowest and the highest utilities achieved by the peers. The percentage increase is the increase of the highest utility in relation to the lowest one. The graph can be interpreted in the following way: on the *x-axis* the percentage of peers is represented and on the *y-axis* we represent the percentage increase, this way we can read that $x\%$ of peers have obtained a gain of at least $y\%$. For example, we can note that 10% of peers have achieved a gain of approximately at least 12%.

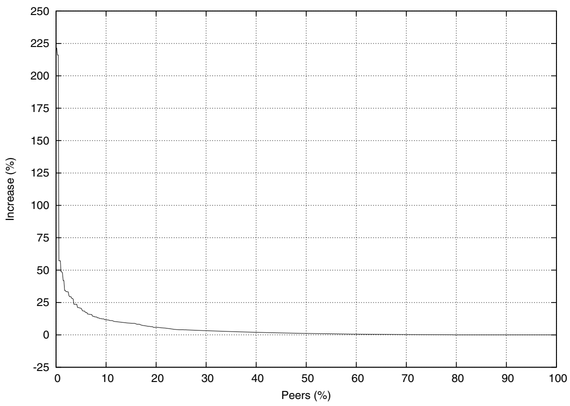

Figure 1. Percentage of peers by utility increasing.

By the figure we can note that a few number of peers increase significantly their utilities changing the services offered to the grid (some of them achieving more than 100% of increase).

We believe that the reason why most peers did not show a meaningful increase is due to the small amount of services in the system. With 6 services and a typical favor of 2 services, at most 4 services can be varied for each peer according to the peer's budget, achieving at most 16 combinations of services. The amount of combinations is given by $\sum_{k=0}^{x} \binom{x}{k}$, where x is the amount of services that can be varied. Since this equation has an exponential increase, if a greater amount of services is considered it might be possible to achieve more peers with a higher percentage increase.

V. Conclusion

In peer-to-peer grid environments, multiple numbers of services can be shared by the peers. Given this fact, we motivated our work from the question whether the choice of services shared with the grid has an impact on the utility obtained by peers. Based on simulation results, we conclude that the choice of services to provide does have a significant impact on the utility obtained for a few peers. Moreover, we believe that the impact of choice will be more important as the number of offered services in the system increases.

We can now use this result to motivate other works. As a future contribution of this work, we intend to evaluate the relationship between the amount of services and the variations in utility which can be obtained by the peers, considering the possible choices of services to offer to the grid. After that, we intend to develop heuristics which choose the best services to offer in a multi-services peer-to-peer grid. We define best services as the ones that yield a greater utility for the peer. In a system where every peer makes use of such a heuristic, we aim at achieving a system equilibrium in which all peers are increasing their utility.

Acknowledgment

This work has been developed in collaboration with HP Brazil R&D.

References

[1] A. A. Selcuk, E. Uzun, and M. R. Pariente, "A reputation-based trust management system for p2p networks," in *CCGRID '04: Proceedings of the 2004 IEEE International Symposium on Cluster Computing and the Grid*, Washington, DC, USA, 2004, pp. 251–258.

[2] N. Andrade, F. Brasileiro, M. Mowbray, and W. Cirne, "A Reciprocation-based Economy for Multiple Services in a Computational Grid," *to appear in Market Oriented Grid and Utility Computing book*, 2008.

[3] R. Buyya and S. Vazhkudai, "Compute power market: Towards a market-oriented grid," in *CCGRID '01: Proceedings of the 1st International Symposium on Cluster Computing and the Grid*. Washington, DC, USA: IEEE Computer Society, 2001, p. 574.

[4] K. Lai, B. Huberman, and L. Fine, "Tycoon: A Distributed Market-based Resource Allocation System," *Arxiv preprint cs.DC/0404013*, 2004.

[5] D. Abramson, R. Buyya, and J. Giddy, "A computational economy for grid computing and its implementation in the Nimrod-G resource broker," *Future Generation Computer Systems*, vol. 18, no. 8, pp. 1061–1074, 2002.

[6] N. Andrade, F. Brasileiro, W. Cirne, and M. Mowbray, "Automatic grid assembly by promoting collaboration in peer-to-peer grids," *J. Parallel Distrib. Comput.*, vol. 67, no. 8, pp. 957–966, 2007.

[7] Y. Benkler, "Sharing Nicely: On Shareable Goods and the Emergence of Sharing as a Modality of Economic Production." *Yale Law Journal*, vol. 114, no. 2, pp. 273–359, 2004.

[8] C. Grothoff, "An Excess-Based Economic Model for Resource Allocation in Peer-to-Peer Networks," *Wirtschaftsinformatik*, vol. 45, no. 3, pp. 285–292, 2003.

[9] B. Cohen, "Incentives Build Robustness in BitTorrent," *Workshop on Economics of Peer-to-Peer Systems*, vol. 6, 2003.

[10] M. Mowbray, F. Brasileiro, N. Andrade, J. Santana, and W. Cirne, "A reciprocation-based economy for multiple services in peer-to-peer grids," in *P2P'06: Proceedings of the Sixth IEEE International Conference on Peer-to-Peer Computing*, Washington, DC, USA, 2006, pp. 193–202.

[11] J. Chase, D. Irwin, L. Grit, J. Moore, and S. Sprenkle, "Dynamic virtual clusters in a grid site manager," *12th International Symposium on High-Performance Distributed Computing (HPDC-12 2003), 22-24 June 2003, Seattle, WA, USA*, pp. 90–100, 2003.

[12] A. Satsiou and L. Tassiulas, "A trust-based exchange framework for multiple services in p2p systems," *Peer-to-Peer Computing, 2007. P2P 2007. Seventh IEEE International Conference on, Galway, Ireland*, 2007.