

Exploring the Catallactic Coordination Approach for Peer-to-Peer systems*

Oscar Ardaiz¹, Pau Artigas¹, Torsten Eymann², Felix Freitag¹, Roc Messeguer¹,
Leandro Navarro¹, Michael Reinicke²

¹ Computer Architecture Department, Polytechnic University of Catalonia, Spain
{oardaiz,partigas,felix,meseguer,leandro}@ac.upc.es

² Institute for Computer Science and Social Studies
Albert-Ludwigs-University Freiburg, Germany
{reinicke,eymann}@iig.uni-freiburg.de

Abstract. Efficient discovery and resource allocation is one of the challenges of current Peer-to-Peer systems. In centralized approaches, the user requests can be matched to the fastest, cheapest or most available resource. This approach, however, shows scalability limits. In this paper, we explore the catallactic coordination as a decentralized economic approach for resource allocation in peer-to-peer networks. The economic model of the catallaxy is based on the self-interested maximization of utility and the negotiation of prices between agents. We evaluate the feasibility of our approach by means of simulations and compare the proposed system with a centralized baseline approach. Our results indicate that while in the catallactic approach the number of control messages exchanged between the peers grows due to the negotiation process, its service provision rate is fairly constant in different dynamic environments.

1 Introduction

Peer-to-peer (P2P) systems are a class of distributed systems or applications to achieve a certain functionality in a decentralized manner. In this model, peers give resources and receive resources in turn. These resources can consist of computing power, storage of data or content, network bandwidth or presence. Application domains are distributed computing, data and content sharing, and collaborative applications.

A number of successful applications like Napster [9], Gnutella [5], and Freenet [4] for file and content sharing, and SETI@home [11] for distributed computing have demonstrated the feasibility of this approach.

* This work was supported in part by the Ministry of Science and Technology of Spain under Contract TIC2002-04258-C03-01 and TIC2001-5193-E, and the European Union under Contract IST-2001-34030 CATNET.

Current P2P systems, however, have mainly focused on the exchange of objects like files and music clips, which are “small”. In future systems, however, the content will be of any form, including audio, video, and large data sets. Peer-to-peer systems may be used to set up multi-cast services for large-scale global audiences, provide services for storing ultra large data sets, and to allow the execution of parallel applications requiring teraflops of processing power. To achieve an efficient performance of such systems, more intelligent decisions than we have in today’s systems are required, concerning particularly from where the content should be retrieved and on which path it should travel.

In this paper, we propose the usage of the economic paradigm of the catallaxis [3] for the decentralized resource allocation in P2P networks. In the catallactic coordination model the decisions of the peers are based on economic principles, being aware that resources like bandwidth, processing power and storage are limited. Peers are in negotiation with each other to optimize their own benefits, and the decision of service provision takes into account their cost and benefits involved.

Recent research in Grid computing has also recognized the value of price generation and negotiation, and in general the use of economic models for trading resources and services and for the regulation of supply and demand of resources in increasingly large-scale and complex Grid environment. Examples are the Nimrod/G Resource Broker and the GridBus project [1; 6].

In order to study the catallactic coordination in P2P networks, we use the CATNET application layer simulator [2], which builds on top of a TCP/IP network simulator the agents to form a peer-to-peer network. We evaluate the proposed system with several experiments and compare the achieved service provision with a baseline system.

In the following section 2 we describe the motivation of our approach. In section 3 we explain the used simulator and the experimental framework. Section 4 contains the evaluation of the proposed system and the discussion of the results. In section 5 we conclude the paper.

2 Decentralized economic coordination with the catallaxy paradigm

Application layer networks such as peer-to-peer networks are software architectures which allow the provision of services by connecting a large number of individual computers for information search, content download, parallel processing or data storage.

In order to keep such a network operational, service control and resource allocation mechanisms are required. Performing such task with a centralized coordinator, however, has several difficulties: 1) A continuously updating mechanism would be needed to reflect the changes in the service demands and node connectivity; and 2) Long latencies for obtaining updates about the nodes at the edge of the network as the diameter of the network grows.

These drawbacks motivate the evaluation of a decentralized coordination concept, which is able to allocate services and resources without having a dedicated and centralized coordinator instance.

The catalaxy coordination approach [3; 7] is a coordination mechanism for information systems consisting of autonomous network elements, which is based on constant negotiation and price signaling. Autonomous agents are able to adapt their heuristic strategies using machine learning mechanisms. This constant revision of prices leads to an evolution of the agent strategies, a stabilization of prices throughout the system and self-regulating coordination patterns [3]. The resulting patterns are comparable to those witnessed in human market negotiation experiments [10].

3 Experimental framework

3.1 The simulator for P2P networks

In order to evaluate the behavior of a P2P system with the catalactic coordination mechanism, we have used the CATNET network simulator [2]. CATNET is a simulator for an application layer network, which allows creating different types of agents to form a network. This simulator is implemented on top of the JavaSim network simulator [8]. JavaSim simulates a general TCP/IP network and provides substantial support for simulating real network topologies and application layer services, i.e. data and control messages among application network instances.

We have implemented two main control mechanisms for the network coordination: the baseline and the catalactic control mechanism. The baseline mechanism computes the service/resource allocation decision in a centralized instance. In the catalactic mechanism, autonomous agents take their decisions in a decentralized way, having only local information about the environment. Each agent disposes of a strategy to take decisions, which targets to increase the agent's own benefit.

In the simulations, we consider a service as the functionality, which is exchanged among the peers in the network. The concept of service and the functions, or "personalities", a peer can assume in the CATNET simulator, are the following:

- **Service:** a service encapsulates a general function performed in the P2P network. A service is the provision of a resource such as computing power, data storage, content, or bandwidth. The service provision includes the search for a resource and its reservation for availability.
- **Client:** A peer may act as a client or consumer of a service. As such it needs to access the service, use it for a defined time period, and then continues with its own program sequence.
- **Resource:** A peer, which is the owner of a required functionality. This functionality, for instance, may represent content, storage or processing power. The functionality, which is required by the clients or consuming peers, is encapsulated in a service.
- **Service copy:** A peer acting as a service copy offers a service as an intermediary, however it is not the owner of the components to provide the service.

It must cooperate with the resource to be able to provide the service. Service copies offer the service to requesting clients.

3.2 Money and message flow

In the simulator, the network activity is characterized by a continuous exchange of control messages and service provision. Different control messages in the two coordination mechanisms are used to accomplish the negotiation between peers.

In Figure 1 we show the money and message flow used in the catalytic coordinated system. The requests from clients are broadcasted and forwarded to the service copies. Compared with the flooded requests model used in Gnutella [5], however, the numbers of hops a request can be forwarded in the simulator is limited. Service copies initiate negotiations with the resources they know to provide the service. Upon successful negotiations the service copies offer the service to the client. If the client accepts, then the service copy provides the service by means of a resource.

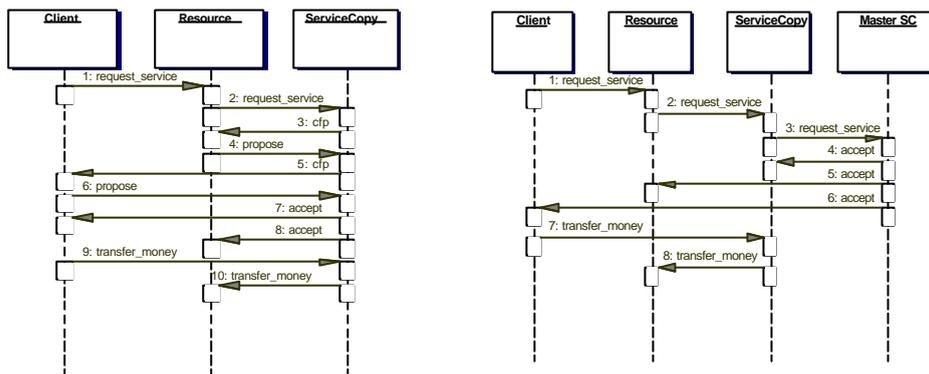


Figure 1. Money and message flows. **Figure 2.** Money and message flows: Baseline approach. Catalytic coordinated system.

In the centralized baseline system (Figure 2), the master service copy (MSC) receives the client requests with additional information through the resource/service copy pairs. Taking into account the distance and availability, it selects a resource/service copy pair and sends back an accept/reject message to the client. The resource allocates the required resource units and the service copy provides the service to the client.

4 Experimental evaluation

4.1 Experimental setup

With experiments we wish to measure if a P2P network coordinated by the catalaxy paradigm is able to successfully provide service to requesting clients. A second goal is to compare qualitatively the obtained results with the centrally coordinated baseline system.

In our experiments we explore as design space of the system the node density and node dynamics of the network (Figure 3). First, we simulate the P2P network with different densities of the service and resource providing agents in a high dynamics environment (Figure 1 experiments 1A-C). Then, we simulate the high node density network in environments with different values of the dynamics (Figure 3 experiments 2A-C).

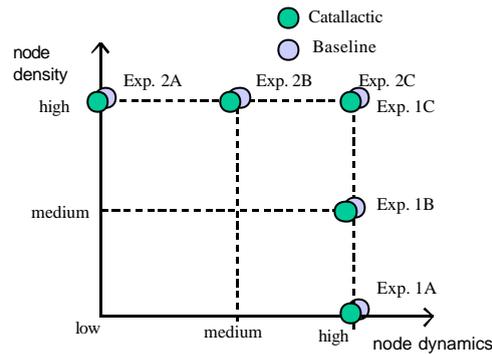


Figure 3. Design space of the system evaluated experimentally. Experiments 1A-C: Node density. Experiments 2A-C: Node dynamics.

In the simulations the input is a trace of client demands with requests for service. The service request specifies the amount of service, a price, and its duration. In all experiments the same demand trace is used.

The physical network topology used in the experiments is organized in three levels of pentagons with leaves on the outer level, such as shown in Figure 4. Although other specific or random topologies of the nodes could be used as well, we applied this topology since it facilitates controlled experiments.

On the top of the physical network an application layer network is build. On each node of the network, peers can be instantiated. Peers are instantiated having one of the previously described types of personalities, which can be a client, service copy or resource agent. Depending on the particular experiment, a node may contain several agents or no agent at all. In the second case, the node acts as a router.

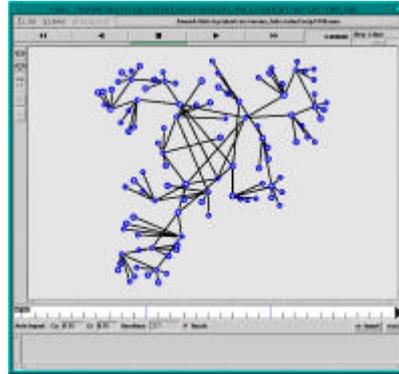


Figure 4. Example of the network topology with approx. 100 nodes used in the experiments.

The relation of the experimental configuration to real world P2P systems is the following: A high value for the dynamics is interpreted to reflect the high level of connection and disconnection in P2P networks. A high level of node density represents the large number of machines with limited capability as found in P2P networks. In the simulations with high node density, we reduce the capacity of the resource agents, in order to represent small machines at the edge of the network. In the low node density scenario, on the other hand, the capacity of the service copy is increased such that the total amount of service available by the network is equal over all experiments.

In Table 1, the configuration of the experiments is detailed.

Table 1. Experiment description.

Input trace	- 2000 service requests generated randomly by 75 clients over a time interval of 100 s. - each request is for 2 service units. - each service has a duration of 5 s.
Node topology.	- 106 physical nodes
Node density	- 75 clients on the leaves of the physical network - different density of resource and service copy agents. Each Resource has one service copy associated. Exp 1A: low node density: 5 resources with capacity 60. Exp 1B: medium node density: 25 resources with capacity 12. Exp 1C: high node density: 75 resources with capacity 4.
Node dynamics	Dynamic behavior: On average 70% of the service copies are connected. Exp 2A: Service copies do not change its state (static network) Exp 2B: Each 200 ms every service copy can change its state (connected/disconnected) with a probability of 0.2. Exp 2C: Each 200 ms every service copy can change its state (connected/disconnected) with a probability of 0.4.

The main parameters we are interested to measure are the number of client requests the network is able to provide a service for in the different scenarios. The scenario we are particularly interested in is the one with high node dynamics and high node density, as this configuration can be related to the conditions found in P2P networks.

4.2 Experimental results

In Figure 5 the service provision rate of a network with *different node density* in a highly dynamic environment is shown (experiments 1A – C). It can be observed that the network using catalactic coordination achieves a higher service provision rates than the baseline system with a smooth decrease for increasing node density. In Figure 6 the service provision rate of a network with high node density in *different dynamic environment* is shown (experiments 2A – C). It can be observed that the service provision rate of the catalactic system is rather independent to the dynamics. The baseline system, on the other hand, decreases with increasing dynamics.

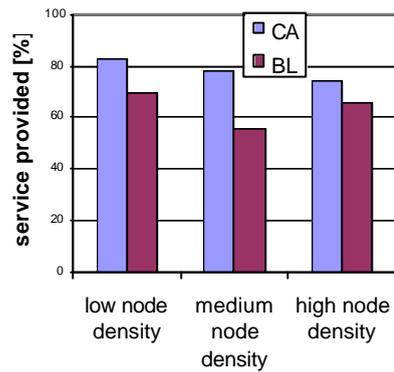


Figure 5. Service provision in % in a highly dynamic network environment with different node density. CA = catalactic coordinated system. BL = baseline system.

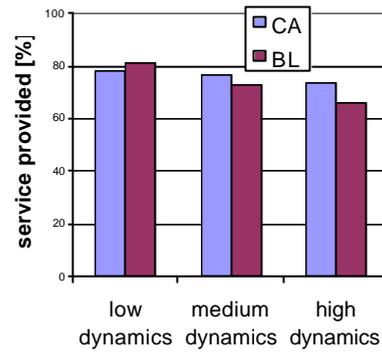


Figure 6. Service provision in % in a high node density environment with different dynamics. CA = catalactic coordinated system. BL = baseline system.

Considering the achieved service provision rate, our experimental results indicate that service provision in networks with many small nodes in a highly dynamics environment could be coordinated successfully by the catalaxy paradigm.

Exploring additional parameters of the system has the potential to provide more insight in the behavior of such a complex system. Currently, we examine the influence of other parameters on the performance tendencies discovered so far. One of the drawbacks we found of the catalactic approach is the time needed to establish a service

provision, which is high due to the negotiation protocol carried out by agents. Other parameters we study is how scale affects the performance of the system.

5 Conclusions

We have first indicated the need for an intelligent decision mechanism for service provision in future P2P networks, which shall cost-consciously decide from where the content should be retrieved and on which path it should travel. Such mechanism should not only achieve the functionality, but also reduce the overall cost to provide this functionality.

We have proposed the catalactic coordination as a decentralized economic approach for resource allocation in P2P networks. In this approach, the decisions of the peers are based on economic principles, being aware that resources like bandwidth, processing power and storage are limited. Autonomous peers negotiate with each other for service provision in order to optimize their own benefits.

With simulations we have investigated if service provision in P2P networks can be achieved by the catalactic coordination approach. We compared the obtained results with a centralized baseline approach. We observed in the experiments that the service provision in the catalactic coordination is rather independent of the dynamics of the network. The service provision capability of the baseline approach appears to be sensitive to the dynamics, reducing its performance in highly dynamic environments.

References

1. Buyya, R., D. Abramson, and J. Giddy, "A Case for Economy Grid Architecture for Service-Oriented Grid Computing". Proc. 10th IEEE International Heterogeneous Computing Workshop (HCW 2001). San Francisco, 2001.
2. CatNet project. "CATNET". <http://research.ac.upc.es/catnet>
3. Eymann, T. and B. Padovan, "The Catalaxy as a new Paradigm for the Design of Information Systems". Proceedings of The World Computer Congress 2000 of the International Federation for Information Processing, 2000.
4. Freenet. 2003. The Freenet home page. <http://www.freenetproject.org>
5. Gnutella. 2003. The Gnutella home page. <http://www.gnutella.com>
6. Grid Computing and Distributed Systems (GRIDS) Laboratory. "GRIDBUS Project". The University of Melbourne, Australia. <http://www.gridbus.org/>, 2002-11-28.
7. Hayek, F.A., W.W. Bartley, P.G. Klein, and B. Caldwell. The collected works of F.A. Hayek. University of Chicago Press, Chicago, 1989.
8. JavaSim Project. "JavaSim". Ohio State University EEng Dept. <http://www.javasim.org/>, 2 A.D.-11-29.
9. Napster. 2003. The Napster home page. <http://opennap.sourceforge.net/>
10. Pruitt, D.G. Negotiation Behavior. Academic Press, New York, 1981.
11. SETI@HOME. <http://setiathome.ssl.berkeley.edu/>