

Streitberger, Werner; Eymann, Torsten

Working Paper

CATNETS final activity report

Bayreuther Arbeitspapiere zur Wirtschaftsinformatik, No. 29

Provided in Cooperation with:

University of Bayreuth, Chair of Information Systems Management

Suggested Citation: Streitberger, Werner; Eymann, Torsten (2007) : CATNETS final activity report, Bayreuther Arbeitspapiere zur Wirtschaftsinformatik, No. 29, Universität Bayreuth, Lehrstuhl für Wirtschaftsinformatik, Bayreuth, <https://nbn-resolving.de/urn:nbn:de:bvb:703-opus-3756>

This Version is available at:

<https://hdl.handle.net/10419/52627>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Lehrstuhl für
Wirtschaftsinformatik
Information Systems
Management

No. 29

2007

Bayreuther Arbeitspapiere zur Wirtschaftsinformatik

Werner Streitberger, Torsten Eymann (University of Bayreuth)

CATNETS Final Activity Report

Bayreuth Reports on Information Systems Management



**UNIVERSITÄT
BAYREUTH**

ISSN 1864-9300

Die Arbeitspapiere des Lehrstuhls für Wirtschaftsinformatik dienen der Darstellung vorläufiger Ergebnisse, die i. d. R. noch für spätere Veröffentlichungen überarbeitet werden. Die Autoren sind deshalb für kritische Hinweise dankbar.

Alle Rechte vorbehalten. Insbesondere die der Übersetzung, des Nachdruckes, des Vortrags, der Entnahme von Abbildungen und Tabellen – auch bei nur auszugsweiser Verwertung.

Authors:

Werner Streitberger, Torsten Eymann (University of Bayreuth)

The Bayreuth Reports on Information Systems Management comprise preliminary results which will usually be revised for subsequent publications. Critical comments would be appreciated by the authors.

All rights reserved. No part of this report may be reproduced by any means, or translated.

**Information Systems and Management
Working Paper Series**

Edited by:

Prof. Dr. Torsten Eymann

Managing Assistant and Contact:

Raimund Matros
Universität Bayreuth
Lehrstuhl für Wirtschaftsinformatik (BWL VII)
Prof. Dr. Torsten Eymann
Universitätsstrasse 30
95447 Bayreuth
Germany

Email: raimund.matros@uni-bayreuth.de

ISSN 1864-9300



IST-FP6-003769 CATNETS

D5.7b Final Activity Report

Contractual Date of Delivery to the CEC:	31 August 2007
Actual Date of Delivery to the CEC:	31 October 2007
Author(s):	Werner Streitberger, Torsten Eymann
Work package:	WP5
Security:	public
Nature:	final
Version:	1.0
Total number of pages:	43

Abstract: This report is the final Activity Report of the CATNETS project. It summarizes the project objectives, the contractors involved, the work performed and the end results. A description of the dissemination activities concludes the report.

CATNETS Consortium

This document is part of a research project partially funded by the IST Programme of the Commission of the European Communities as project number IST-FP6-003769. The partners in this project are: LS Wirtschaftsinformatik (BWL VII) / University of Bayreuth (coordinator, Germany), Arquitectura de Computadors / Universitat Politècnica de Catalunya (Spain), Information Management and Systems / University of Karlsruhe (TH) (Germany), Dipartimento di Economia / Università delle merci Ancona (Italy), School of Computer Science and the Welsh eScience Centre / University of Cardiff (United Kingdom), Automated Reasoning Systems Division / ITC-irst Trento (Italy) / Chair of Business Administration and Information Systems – E-Business and E-Government / University of Mannheim (Germany)

University of Bayreuth

LS Wirtschaftsinformatik (BWLVII)
95440 Bayreuth
Germany
Tel: +49 921 55-2807, Fax: +49 921 55-2816
Contactperson: Torsten Eymann
E-mail: catnets@uni-bayreuth.de

Universitat Politècnica de Catalunya

Arquitectura de Computadors
Jordi Girona, 1-3
08034 Barcelona
Spain
Tel: +34 93 4016882, Fax: +34 93 4017055
Contactperson: Felix Freitag
E-mail: felix@ac.upc.es

University of Karlsruhe

Institute for Information Management and Systems
Englerstr. 14
76131 Karlsruhe
Germany
Tel: +49 721 608 8370, Fax: +49 721 608 8399
Contactperson: Christof Weinhardt
E-mail: weinhardt@iism.uni-karlsruhe.de

Università delle merci Ancona

Dipartimento di Economia
Piazzale Martelli 8
60121 Ancona
Italy
Tel: 39-071- 220.7088 , Fax: +39-071-220.7102
Contactperson: Mauro Gallegati
E-mail: gallegati@dea.unian.it

University of Cardiff

School of Computer Science and the Welsh eScience Centre
University of Cardiff, Wales
Cardiff CF24 3AA, UK
United Kingdom
Tel: +44 (0)2920 875542, Fax: +44 (0)2920 874598
Contactperson: Omer F. Rana
E-mail: o.f.rana@cs.cardiff.ac.uk

ITC-irst Trento

Automated Reasoning Systems Division
Via Sommarive, 18
38050 Povo – Trento
Italy
Tel: +39 0461 314 314, Fax: +39 0461 302 040
Contactperson: Floriano Zini
E-mail: zini@itc.it

University of Mannheim

Chair of Business Administration and Information Systems – E-Business and E-Government
L9, 1-2
68131 Mannheim
Germany
Tel: +49 621 181 3321,
Fax +49 621 181 3310
Contactperson: Daniel Veit
E-mail: veit@uni-mannheim.de

Changes

Ver- sion	Date	Author	Changes
0.1	08.10.2007	WS	First version
0.2	15.10.2007	WS	Small corrections
1.0	15.10.2007	WS	Final release

Content

- 1 Project Objectives..... 5
 - 1.1 Background 5
 - 1.2 Objectives..... 5
 - 1.3 Approach 6
- 2 Activities performed and results achieved 7
 - 2.1 Theoretical and Computational Basis 7
 - 2.2 Simulator 9
 - 2.3 Prototype 11
 - 2.3.1 Architecture11
 - 2.3.2 Catallactic-enabled applications.....15
 - 2.3.3 Standards16
 - 2.3.4 Implementation17
 - 2.3.5 Evaluation of prototype performance17
 - 2.3.6 Conclusions18
 - 2.4 Metric evaluation 19
 - 2.4.1 The metrics pyramid19
 - 2.4.2 Technical and economical metrics21
 - 2.4.3 Statements on applying the Catallaxy to Application Layer Networks...24
 - 2.5 Future Research on Properties of Catallaxy applied to Computer Networks26
 - 2.6 Conclusions 28
- 3 Publishable Exploitable Knowledge and its Use31
- 4 Dissemination and Use.....33
 - 4.1 Dissemination to the Scientific Community..... 33
 - 4.1.1 Organization of Scientific Meetings.....33
 - 4.1.2 Refereed Journal Publications34
 - 4.1.3 Publications & Attendance at Scientific Meetings35
 - 4.1.4 Participation in Scientific Networks39
 - 4.2 Dissemination to a General Audience 40
 - 4.2.1 The Website.....41
 - 4.3 Public deliverables..... 41
- 5 Co-ordinator Contact Details43

1 Project Objectives

1.1 Background

Future Grid network technology will face the problem of the efficient provisioning of services to clients by a scalable and dynamic resource allocation (matching) mechanism. The objective of CATNETS is to determine the applicability of a decentralized economic self-organization mechanism for resource allocation in application layer networks (ALN), by

- (1) producing a 'proof-of-concept' prototype in a real ALN and by
- (2) evaluating its performance against existing resource brokerage approaches in a simulated ALN.

The term ALN integrates different Internet overlay network approaches, like Grid and P2P systems. The allocation of resources in these networks (e.g. matching of demand and supply, deployment of service instances, and service discovery) can principally be conducted in a centralized (e.g. using resource brokers) or in a decentralized fashion (e.g. using self-organizing mechanisms). Centralized approaches reach their limits with increasing network size and growing numbers of elements; self-organizing approaches thus gain attention, e.g. in IBM's Autonomic Computing initiative.

The CATNETS project investigates a 'free market' economic self-organization approach, the 'Catallaxy' by Friedrich A. von Hayek, as the basis for self-organizing resource allocation in ALNs. A preliminary evaluation of 'Catallactic' mechanisms in the FET assessment project CATNET (IST-2001-34030) by simulation has shown positive results, upon which the CATNETS project builds. The performance measurements of the 'proof-of-concept' implementation are compared with the simulation results, with the goal of being able to make a substantiated statement on the applicability of economic self-organization as a major component of ALN networks. A positive evaluation of the Catallactic approach would have a high potential impact, with new possibilities for resource brokering in future ALN, and maybe for self-organization in computing in general.

1.2 Objectives

Application-layer networks (ALN) are envisioned as large, complex computer networks that allow the provisioning of services requiring a huge amount of resources. They connect large numbers of individual computers for information search, content download, parallel processing or data storage. Common concepts are Grid computing (used for distributed processing) and Peer-to-Peer-(P2P)-Computing (used for distributed data storage and access). Allocating and scheduling the usage of computing resources in ALNs is still an open and challenging problem.

This project proposes an innovative research for realizing resource allocation in ALNs by applying the economic self-organization and decentralized paradigm, F.A. von Hayek's „Catallaxy”, which goes well beyond the state-of-the-art of existing centralized economic mechanisms in Grid or ordinal mechanisms in P2P networks.

This project is based on the preceding CATNET assessment project in the EU's FET Open programme, which has shown promising results of using this „free market” mechanism in highly dynamic and large computer networks.

The main objective of the project is to provide a significant statement about using "free market" (catallactic) mechanisms in application layer networks.

Our conclusions concern their applicability and implementation possibilities in Grid/P2P middleware, providing a prototype, and performance results and further insight in their behaviour, both from a technical and economic point of view, using simulations.

1.3 Approach

This objective will be achieved through a twofold approach. On the one hand we will explore the feasibility of Catallaxy for real ALNs by developing a prototype. On the other hand we will investigate the potential and limitations of the Catallaxy in ALNs by means of simulations. This approach leads to the following tasks to advance in the knowledge of both the practical and theoretical aspects of Catallaxy in application networks:

- **Proof of Concept:** The project investigates how „Catallactic” mechanisms can be implemented for resource allocation in real application layer networks. This requires a specification of the components that use „Catallactic” mechanisms and how these components can be integrated in the middleware of current Grid and P2P platforms. Their functionality covers resource brokerage, resource discovery, and re-deployment of services. The prototype represents an important result to demonstrate the feasibility and applicability of the approach.
- **Evaluation by simulation:** The project extends the simulations used in the assessment project, to compare the performance of real application layer networks using „Catallactic” mechanisms with that of using other economic mechanisms by conducting simulation experiments. While the implementation of the components is the same for the „proof of concept” prototype and the simulation modules, the simulation allows investigating in more detail the aspects of e.g. scalability, topology influence, or connection reliability. The experience made during the implementation and the performance of the Catallactic enhanced simulation as well provides important feedback for the design of real systems.
- **The challenge:** To achieve our objective the project faces the challenge combining contributions both from computer science and economics to address the features of coming Grid and P2P applications and infrastructures. Catallaxy has been proposed as a model to describe the behaviour of complex and large scale real world economies. However, its results have not yet been transferred to coordinate large and dynamic computer networks. In the previous assessment project CATNET we have conducted an initial study on the potential of the approach. CATNETS extends this initial step to evaluate this approach as an economic mechanism to regulate the dynamic and decentralized services required in the emergent connected society using the next generation Grid and P2P systems.

2 Activities performed and results achieved

2.1 Theoretical and Computational Basis

What makes economics so attractive for computing environments is that its central research question lies in the efficient allocation of resources, provided by suppliers and in demand by customers. In computing environments like Grid Computing, the resources in question are processor time or storage space, while economic actors are represented by software agents or web services. It appears that, by just implementing markets in computing environments, the satisfying ability of economics might be viable for creating cost-effective computer architectures. However, between the mostly descriptive economic concept and the normative technical implementation lies a fundamental gap, requiring selective choice of how actors, resources, goods, and markets are modeled and embedded into a technical environment.

In Application Layer Networks, participants offer and request application services and computing resources of different complexity and value - creating interdependent markets. In CATNETS, these complex interdependencies are broken down into two types of interrelated markets:

- (1) a resource market - which involves trading of computational and data resources, such as processors, memory, etc, and
- (2) a service market - which involves trading of application services.

This distinction between resource and service is necessary to allow different instances of the same service to be hosted on different resources. It also enables a given service to be priced based on the particular resource capabilities that are being made available by some hosting environment.

The scenario that is envisioned in CATNETS is, that there is a set of basic services (e.g. services to create a PDF or to convert a MP3 file), a set of complex services demanding these services for a specific job (e.g. an application wants to create a PDF file), and a set of resource services capable providing computational resources for executing these services (e.g. a processor, main memory, and a hard disk for creating the PDF file). However, an agent that is requesting a service is unaware of the resources the requested service requires to be carried out.




Figure 1: CATNETS Scenario – Service Market and Resource Market

Figure 1 illustrates the CATNETS scenario. A complex service is requesting a PDF creator service, which will be allocated to the agent. Furthermore, the required resources (CPU and hard disk) are allocated to the service. The service acts as a trading intermediary, i.e. the service knows what the agents are demanding and which resources are available for executing the services.

The core contribution of the CATNETS project is the quantitative comparison between common centralized economic allocation mechanisms and decentralized negotiation formats based on von Hayek's Catallaxy. Therefore, several metrics have been defined in order to identify the quality of the economic allocations. The results show, that the applicability of allocation mechanisms highly depends on the parameterization of the individual setup. Core issues along which the identification of the appropriate mechanism has to be aligned at are:

- the size of the allocation problem,
- the communication intensity,
- the distribution of the prices offered by the participants and
- the dynamics of the market.

All these parameters again depend upon the industry branch in which the individual application, for which the mechanism should be deployed, is located. The approaches investigated within the CATNETS project may be subsumed in the field of Grid Economics. In this field, currently strong efforts are bundled in order to identify methodologies that are applicable for dynamically allocating computational resources to applications. The vision of this field is to enable a seamless integration of distributed hardware for the computation of heterogeneous front-end applications. The idea is to allow for dynamic allocation of resources in order to determine the prices of computational resources along the time. In order to lay the fundamentals for such architecture, substantial efforts have to be carried out in the Grid middleware domain. Applying the results from CATNETS anticipates a fully functional Grid middleware, which is capable of ex-ante determining the time specific jobs are running. This implies a component, which judges the runtime of jobs stemming from heterogeneous applications.

Having such a component in place will allow the application of market based allocation schemes such as they have been proposed in this work.

The key ideas of our approaches have been presented in different communities. A large number of experts from the e-Infrastructure community, the Grid community as well as the SOA and the distributed systems communities see great potential in these approaches. Several in depth co-operations have been started here.

2.2 Simulator

The CATNETS simulation environment poses low requirements on simulation runs. Implemented in pure JAVA, the CATNETS simulator runs on all machines which are supported by the Java Runtime Environment. This enables small scale simulation on desktop machines and large scale simulation on multi-core 32 and 64 bit server machines. The pure JAVA implementation gives the possibility of easy adoption of the current code to new simulation scenarios of application layer networks in utility computing or autonomic computing areas.

The high resource and service abstraction and the two-tiered market implementation support various areas of application. The introduction of resource bundles enables the modelling of not only a specific resource type like data resources, but complex resource products for future peer-to-peer enabled Grid applications. We assume, visualization techniques and local resource managers are in place and offer an abstract resource bundle to services. In CATNETS, we implemented two different allocation policies for these resource bundles, a centralized auctioneer using a multi-attributive combinatorial auction and the catallactic approach using bilateral bargaining.

A resource provider can select between two implemented resource models for his resource service, a shared and a dedicated resource model. Using the dedicated resource model, super computing or autonomic computing can be simulated, whereas the shared resource model represents scenarios of the utility computing field. Both resource models allow co-allocation of resource bundles from different resource providers. This allows the simulations of basic services with high resource demands like the execution of batch jobs in the super computing area. Not all possible co-allocation scenarios are supported by the current catallactic implementation because of its high complexity. In the catallactic approach, we assume the co-allocated resources have the same resource bundle size, capacity and product id. The central approach supports all co-allocation combinations which allows allocation of bundles with different size and capacity.

The service market decouples the service requests from the resource market. A complex service doesn't have to know how many resources there are and how many resources he needs for his service. The complex service can focus on creating value added services to the user. Currently, the complex services sequentially request a list of basic service. This could be enhanced in future releases of the CATNETS simulator with more sophisticated workflow engines. As on the resource market, a service allocation policy is applied to allocate services. In CATNETS, we implemented two allocation approaches, a continuous double auction and the catallactic bargaining.

Supporting both allocation mechanisms, the simulator provides proactive and reactive interfaces for software agents. Proactive agents act on their own. They periodically check, if there is new demand or supply and send their bids to the auctioneer. In

the reactive agent model, the agents wait for new events like incoming messages before they act on new situations.

The simulator provides a rich messaging model including a large set of different message types, a P2P messaging layer with flooding, load-link dependent message delay, and a simple message failure model. This enables the simulation of real life influences on the resource allocation approaches. Interfaces ease the implementation of an improved P2P layer or new message types.

Various tools were developed to support the configuration of the simulator. Scenario generators help to create new scenarios or configure the catallactic market. The simulator also supports plain text file based configuration. This allows fast reconfiguration of the scenarios between simulation runs.

A large set of metrics was implemented in the simulator. This set of metrics helps to debug and evaluate the simulation runs. Technical and economic metrics are written into text files which are evaluated with MATLAB scripts ex-post. Again, the use of text files gives the possibility to use any tool for analysis.

The simulation environment and all developed tools will be released under an open source licence. This will give other researchers the possibility to modify and extend the simulator for their own research.

The simulator allows carrying out experiments up to a few thousands of nodes. However, very large-scale experiments with several thousands of nodes consume lots of main memory. The automated scenario generator is tested with a scenario up to 10000 nodes and 100000 agents. This requires a machine with 8GB of main memory. The simulator is able to read a scenario of this size and to start the simulation with 1000000 requests. But, it takes up to weeks to finish such very large simulation runs. Accessorily, the memory management of the Java Virtual Machine has to be optimized by modifying heap size and memory consumption parameters, and the Linux file system parameters have to be changed to be able to handle such large simulation runs. Therefore, no very large scale experiments are evaluated because the computational resources are needed to simulate the smaller experiments for the scenarios of this deliverable.

In the largest scale simulations, we could only apply the Catallactic mechanism, but not the centralized mechanism because components of the centralized auction implementation consume a large number of resources depending on the size of the scenario. We observe that the Catallactic mechanism could principally achieve very good service allocation. However, this mainly depends on configuring its parameters correctly.

The development of the simulator and the experiments show a high complexity of the agent strategy in catallactic system. Again, the performance depends on numerous parameters which are not easily configurable appropriately.

In the simulations, a number of advantages and disadvantages of the decentralized mechanism compared to the centralized catallactic mechanism have been identified:

- **Number of messages:** The centralized mechanism used a significantly smaller number of messages for communication. The bargaining of the catallactic mechanism had a significantly higher cost in terms of messages exchanged. The number of messages in the implemented auctions remains always the same for one allocation, whereas the number of messages varies along the

iterative negotiation steps in the decentralized allocation approach. By limiting the number of negotiation rounds, a lower number of messages can be achieved.

- **Negotiation attributes:** In the catallactic mechanism, only single-attribute negotiations could take place, whereas the auctioneer mechanism of the centralized approach could handle multi-attribute negotiations. In our particular case, this limitation was imposed by the agents' strategy implementation. The learning and decision algorithms were implemented to cover only single-attribute negotiations. An extension to multi-attribute, however, has several difficulties, like matchmaking of multi-attribute services in a decentralized search.
- **Negotiations:** The implementation of the agents was based on a model in which each Catallactic agent could only be involved in one negotiation at a time. This model poses serious limitations on performance. A more efficient model with parallel negotiations isn't implemented due to its additional complexity. The same limitation holds with the centralized allocation approach. Parallel bidding is not supported.

2.3 Prototype

2.3.1 Architecture

Current Grid Computing architectures exhibit fairly static resource infrastructure which are connected by physically stable links. The shift to a pervasive Grid, that could exist ubiquitously, demands for a more dynamic consideration of resources and connections. The market is understood to be a decentralized control mechanism for services and resources. Figure 2 shows this service oriented model.

A complex service could be represented by a proxy which needs (remote) basic service capabilities for execution – with support for a service selector instance. Complex services are therefore shielded from the details of the resource layer. A basic service is split into the basic service logic and a resource allocator. The logic is able to negotiate with the complex service and to translate the requirements for service execution on a resource instance (e.g. CPU, and storage, etc.). A resource allocator gets the resource specification and broadcasts the respective demand to the local resource managers. This comprises bundles and collocative negotiations. Bundles are understood as an n-tuple of resource types (e.g. a 3-tuple would be: CPU, storage, and bandwidth); co-allocation describes obtaining resources for one single service transaction from various local resource managers simultaneously. Local resource/job scheduling is not the focus of the project. It is expected that a local resource manager hides all details of the allocation.

On the first market, complex service and basic service negotiate; an agent managing a complex service acts as a buyer, the basic service agent as a seller. The same market roles can be found at the resource layer, the resource allocator is the buyer agent, the local resource manager acts as a seller agent. Contemplating the second market, it is a "n" to "k" market: "n" basic service copies can bargain with "k" resource services. This takes dynamic resources into account. Resources are in our view enti-

ties that can fail like basic services, and which are subject to maintenance and inspection procedures or link failures.

The Basic Service is a standardized service for query job execution. At the application layer, the Basic Service consists of:

- a seller entity on the service market;
- an entity to translate a query based on resource demands.

The main functionalities of the Resource Co-Allocator at the resource layer are:

- Representing the buyer entity on the resource market.

Entities of a seller on the resource market are able to provide a set of resources via the Local Resource Manager (LRM). The Resource Agents act on behalf of these LRMs, which hide the physical resources behind them.

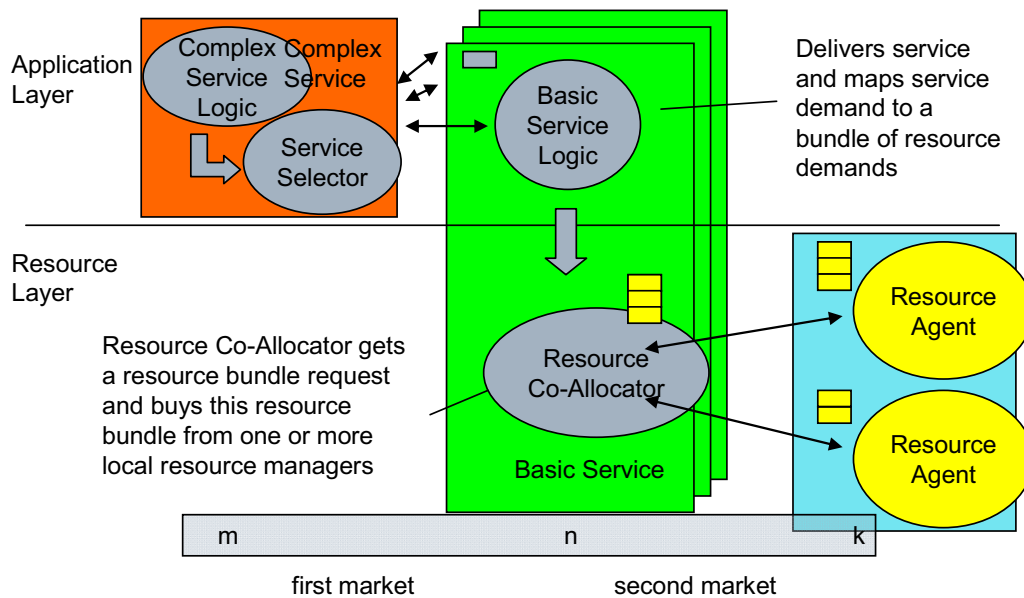


Figure 2: Catallactic Service Oriented Model

We expect an ALN to be built from basic services that can be dynamically combined to form value-added complex services. These basic services require a set of resources, which need to be co-allocated to provide the necessary computing power.

Therefore, the introduction of new services into an ALN, due to the dynamic nature of the environment, precludes any manual or static configuration and demands a self-organization approach. One goal of self-managed network services is to move away from individual system configuration management to policy management. This approach brings a higher level of abstraction to management, by introducing a policy from which the configuration is derived, allowing components of the infrastructure to apply these derived configurations to the individual systems across the environment.

In this context, the resources associated with self-managing services make use of Service Level Agreement-based (SLA) policies for services and resources, mapping required SLAs to resource needs, discovering resources that guarantee an adequate Quality of Service (QoS), allocating resources to ensure that allocation policies are

meet, and providing a management interface to monitor a control service life-cycle. Because of the dynamicity of the environment for which our approach is intended, the service allocation framework must address some specific issues:

- Situation: a service must be aware of its location and the availability of peer services within its “vicinity”, with which to collaborate;
- Dynamic (re)configuration: usage patterns from service users are unpredictable, therefore neither the location nor the number of service instances could be known in advance. New instances must be created and located as needed;
- Topology neutrality: services deployed in the ALN could have very different interaction topologies. Some will be structured in a hierarchical overlay, like content distribution networks, while others interact in a closely connected P2P overlay;
- Autonomy: a service and the resources it uses will span multiple administrative domains, so each resource provider should be allowed to take decisions autonomously.

The economic-based resource allocation model has been implemented in the Grid Market Middleware (GMM), which provides the mechanisms to register, manage, locate and negotiate for services and resources. It allows trading agents to interact with each other based on their requirements/demands and engage in negotiations. Furthermore, the middleware offers a set of generic negotiation mechanisms, on which specialized strategies and policies can be dynamically plugged in. The middleware – as shown in Figure 3 – has a layered architecture, which allows a clear separation of platform specific concerns from the economic mechanisms, to cope with highly heterogeneous environments.




Figure 3: Layered middleware architecture

Figure 4 shows the interaction between the applications and the GMM, which is based on the WS-Agreement specification. The use is initiated from within the prototype, at the application level; the interaction between the application and the middleware is as follows: a user issues an application request, which is interpreted by the prototype and converted into an application service request; then the application determines which Grid service(s) is (are) required to fulfil the specific application service request. These Grid services represent either software services (e.g. query services or data mining services) or computational resources. The application service trans-

lates these requirements into a WS-Agreement template, which is submitted to the GMM.

The GMM searches among the available Grid/Web Services, which have registered their particular service specifications, such as: contractual conditions, policies and QoS levels. When a suitable Grid service provider is found, the application requirements are negotiated within the middleware by agents who act on behalf of the service providers as sellers and the application as buyers. Once an agreement has been reached between the trading agents, a Grid service instance is created and a reference is returned to the application, which then can invoke it.

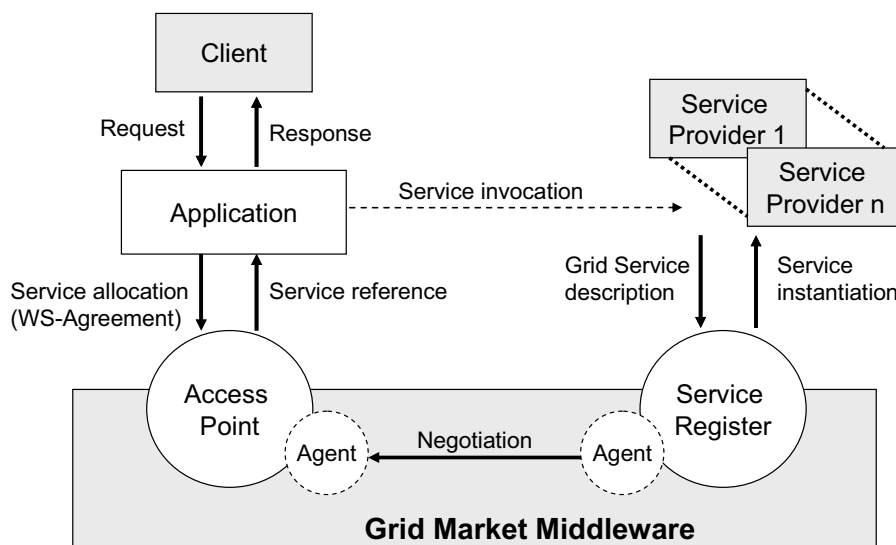


Figure 4: Interaction between the Application and the Grid Market Middleware

To take full advantage of these market mechanisms, it will be necessary to include some logic at the application level by specifying a maximum price for a “Web Service type” the application is interested in. This maximum price is like a budget for the buyer service agent that could be spent to buy basic services and resources. The logic at the application level will include a policy document on which the decision is taken by the application to accept or reject the Web Service instance that has been found on the market. The policy document in this case will take the “price” as the characteristic to be taken into consideration, and this is achieved by checking that this is less than the budget proposed initially by the application.

The benefits of defining this architecture have been the following:

- Following this architecture, the different agent strategies could be implemented by substituting mainly only the components referring to the Economic-Algorithms layer.
- The communication between the application and the middleware has been reduced to the communication to a Catallactic access point, which by means of the WS-Agreement specification transmits all required information to the middleware.

2.3.2 Catallactic-enabled applications

The catallactic paradigm as introduced in predecessor Catnet project has been implemented in a concrete middleware and prototype applications in the scope of CATNETS project. The experience gained on such endeavor has resulted in both deceiving and promising conclusions, shared in equal proportions.

In the deceiving side, important trouble during prototype calibration has come from the requirement of dealing with decentralized decision makers (the trading agents) in a real, networked, infrastructure. Such a development is pioneering, since state of the art approaches to fully decentralized markets based on bargaining agents have been based purely in simulations. In that sense, “touching the ground” of real deployment has proven hard due to the uncertainty and lack of control in such “engineering with complexity” tasks. Emergent engineering has been largely coveted by large distributed systems engineers, but to date just initial steps in form of proposals have been achieved. In our view, and summarizing our experience, more advances through extensive testing in both software lifecycle management and practical deployment tips need to be realized in order to reach maturity.

In the promising side, we have been indeed able to design, implement and deploy a fully decentralized prototype incorporating emergence and self-organization using state of the art tooling. The GMM implementation has been proven as useful in several applications. The use of the Catallactic middleware has been shown by two applications: COVITE being available as Grid Service and Data Mining tools given as Web Services. Both of these applications have been Catallactic-enabled within the project. From the experience obtained, we have found that applications provided as services can be enhanced with reasonable effort to interact with the GMM. Additionally, the GMM is being useful as an infrastructure root for further development in Grid Markets research project, as in the case of SORMA. The open issue with the catallactic-enabled applications is to achieve improved system control, in the form of more predictable outcomes out of the emergent properties of the markets.

From the application point of view, the fact of having participants offer and request for application services and computing resources of different complexity and value in a distributed environment leads to the creation of interdependent markets. In such interrelated markets, allocating resources and services on one market inevitably influences the outcome on the other markets. A common approach of many other Grid market concepts is to allocate resources and services by relying on the presence of centralized resource/service brokers. However, the complex reality could turn such approaches useless, as the underlying problem is computationally demanding and the number of participants in a worldwide distributed environment can be huge.

Different examples of application scenarios can be constructed which benefit from using the Catallactic markets in combination with different auction mechanisms in the Grid. This leads to an advantageous flexibility in terms of fulfilling the requirements and needs of services and resources within the applications and hides all the complexity to the users. Let us consider an application scenario that requires a highly specialized service such as medical simulation service or visualization service, while another application requires a specific mathematical service. The mathematical service is more or less standardized and there are several suppliers offering this service, and an instance of a catallactic market could be initiated and based, for example, on a normal double auction. The medical simulation service, however, does not have

many service suppliers; therefore the liquidity of the market trading such services may be low. In such cases, an instance of a market could be initiated and be based on English auction mechanism. Other types of applications enable creation of Virtual Organizations (VOs) for planning, scheduling, and coordination phases within specific projects or businesses, and allows the users of a VO to interact among them for the duration of VO. The ability of a free-market economy to adjudicate and satisfy the needs of VOs, in terms of services and resources, represent an important feature that markets, through the auction mechanisms, can provide too. Such VOs could require large amount of resources which can be obtained from computing systems connected over simple communication infrastructures such as Internet.

In conclusion, catallactic-enabled applications are well motivated and address a real need of current realistic Grid scenarios. We have developed the first prototype available which is able to deploy the complex catallactic behavior in real Grid applications. The experience gained is valuable as it is, but it can be also profited by engineers in the field of “engineering with complexity and emergence”, where prototype implementations and deployments in real testbed applications are increasingly necessary to advance the state of the art.

2.3.3 Standards

Service Level Agreements (SLAs) provide a contract between an application user requiring services/resources, and application providers determining what should be made available for external use. To enable service/resource sharing/usage in application environments, SLAs may be used to define: (a) requirements that such an application would place on services (and resources) owned by a third party; (b) check whether these requirements have been met during use. An SLA may also specify the penalty that a service provider may incur if terms in the SLA are violated.

Currently, SLAs are defined in a static manner, i.e. the terms within an SLA must adhere to strict constraints, and are monitored during application execution – such as in WS-Agreement. However, within many applications, it is often difficult to define such constraints very precisely, thereby leading to a large number of violations. There is a need to modify an agreement that had already been established, especially if the agreement is used at a time much later than when the agreement had been defined. These requirements relate to comparing the cost of re-establishing a new agreement vs. being able to adapt an agreement that is already in place. Secondly, there is a need to support flexibility in the agreement if an agreement initiator is not fully aware of the operating environment when the agreement is defined. In this case, the agreement initiator may not have enough information to determine what to ask for from a provider. This is likely to be the case when an agreement initiator or provider operates with imprecise knowledge about the other party involved in the agreement.

Specifications which have been applied for the development of the prototype have mainly used the concept of SLA using the WS-Agreement protocol. The specification allowed describing the services needed by the users' application. The protocol for the exchange of WS-Agreement messages between the application and middleware needs to be developed for further negotiation interaction, which has been identified as a limitation.

2.3.4 Implementation

The implementation of the prototype took advantage of the functionalities already provided by available toolkits, like Diet agents, JXTA, GT4 (the middleware selection process as well as the GMMs' early design with them is described in the first year deliverable of WP3 [Del05a]). When running initial experiments with the developed prototype, however, limitations in the practical use of these toolkits have been observed, like a limited number of messages which could be sent with Diet, and the difficulty of JXTA to work correctly with a small number of nodes.

The observed limitations of the Diet toolkit affected the initial design of the performance measuring framework. Another design has been finally implemented which did not rely on the messaging mechanism of Diet, such that in the current prototype we do not have a limitation concerning this issue.

The limitation of JXTA affects the scope of deployment of the prototype in the sense that in small scale scenarios the delivery of messages by JXTA is not reliable. In the environment of the cluster where the prototype has been used, the identified limitation did not appear by making use of particular JXTA messaging services for this context. From the experience obtained, for complete decentralized scenarios, DHT implementations like Pastry could have been a better choice for the implementation of the communication and search.

2.3.5 Evaluation of prototype performance

The prototype has been deployed in a cluster of Linux machines. Several experiments have been made with different type of economic agents, and varying the different parameters of the experiment configuration. The experiments showed the behavior of the measured parameters for concrete experimental settings. The character of the experimental results is rather that of confirming the implementation feasibility in terms of a prototype. The comparison of different agent approaches by means of the developed prototype is difficult. A large number of parameters remain uncontrolled due to the use of a real environment. The difficulty in configuring the Catallactic agents in the sense of being able to chose the most appropriate values for the parameters makes it appearing too early for stable quantitative studies.

The level of the prototype imposed certain constraints on the system compared with what could be a production-quality implementation. As a consequence, quantitative results are only achieved within the scope of the prototype. Our performance results are experiments taken at different stages of the prototype development. Although they revealed the behavior of the system in the current experimental configuration, they gave feedback on the implementation and hints on the complexity of applying Catallaxy in real systems.

The flow of a typical prototype experiment is quite simple, if we consider the large number of automated steps from the original client request until the moment when the end point reference for service execution is returned back. An example of an experiment with 4 nodes is the following: Two nodes host a BS each and the Data Mining Web Service and other two nodes host the CSs, access points and clients. The Web Services are exposed in Tomcat servers. Access for execution of these Web Services on the resource node is what is traded between BSs and CSs. The experi-

ments consist of launching 2 clients concurrently, which use each one of the CS as broker. Each client makes 100 requests to the CS in intervals of 2 seconds. Whenever a CS wins a bid with a BS, it invokes the Data Mining Service in the selected node, and the resource in the corresponding node gets locked for the duration of the service execution. We measure the selling prices of the BSs and observe the proportion of successful CFPs issued by the CSs. The development of the prototype has allowed assessing the feasibility of implementation, providing a flow of execution from the Client to the GMM access point, and from the Complex Services to the Basic Services, in the sequence depicted in Figure 5.

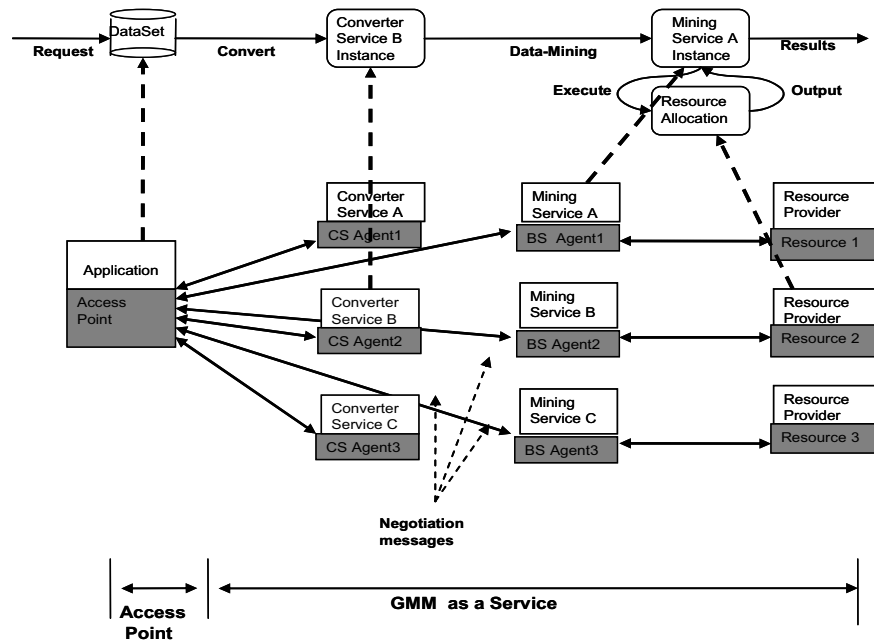


Figure 5: Flow from Client request until Basic Service trading resources.

In general the performance results have been promising in the sense that they show coherent self-organizing behaviours in the bargaining agents. This includes coherent reaction to offer/demand variations, varying computational size of services to be executed and the effect of background loads in the Grid. Improved control of such behaviors could be achieved with further refinements of both multi-agent coordination protocols (i.e. the catallactic agents themselves and a measurement infrastructure incorporating latest advances in complexity management). Another interesting property which could not be assessed in the prototype (given the material limitations of the number of physical nodes available) is that of scalability. However, the results from the simulator on improved scalability of decentralized catallactic agents over centralized approaches offers a good insight on the “nice” scalability properties of the catallactic agents.

2.3.6 Conclusions

The implementation of two prototypes with a Catallactic based mechanism as well as generalisations of application with the GMM have raised interesting conclusions outlined below:

- COVITE prototype supports the creation of Virtual Organisations (VOs) via a centralised system. All resources and services are allocated by a central authority at the creation of the VO. A decentralized solution, as Cat-COVITE prototype presented above, in which services and resources are to be allocated in peak/busy times to application users / processes is a straightforward way to deal with such demanding situations, while there is no other possibility of adding more resources or the availability of services is limited.
- A second solution, a Cat-Data Mining prototype, has been proposed and developed. The main problem data mining services are addressing is of data that is typically too inconsistent and difficult to understand into such forms that are more compact, useful and understandable. This can be achieved via specific data-mining methods for pattern discovery and extraction. This process can be structured into a discovery pipeline/workflow, involving access, integration and analysis of data from disparate sources, and to use data patterns and models generated through intermediate stages. The GMM addresses the need of such data mining services to be found just-in-time and used by the application services.
- The Catallactic markets are expected to provide members of VOs a fair price for services needed by the users' application.
- The Catallactic markets are also expected to provide a fair welfare distribution among the participants, such as users' application, and service and resource providers.
- The Catallactic mechanism is expected to help systems in discovering and selecting resources and services on demand and just in time, as application processes can make use of third parties services or can demand more and more resources.
- Relevant prototype metrics are proposed and a composite index - social utility factor, is being calculated as a single index, so that comparison and evaluation of different scenarios, as well as between prototypes and simulators can be realised.

2.4 Metric evaluation

2.4.1 The metrics pyramid

It is often useful to be able to compare two allocation methods using a single index or number. Such an index provides an aggregated behaviour of an allocation method with reference to a number of features. The application of statistics is necessary to achieve the index, but is useful only in conjunction with a detailed measurement framework. The results of the measurement are intended to allow providing feedback on system behaviour, improve chances of successful adaptation, improve implementation, and increase accountability.

Figure 6 shows the logical structure of data and indices. It sums up the underlying methodology establishing the shape of a pyramid. Data are the basic units of information, collected through technical monitoring of the application layer network. Parameters which are likely to be of significance within the application and the resource allocation mechanism have to be selected for measurement. These parameters define the raw disaggregated data. To ease the analysis of the raw data, they are collected from different experiments (simulation runs) into plain text files. Disaggregated

indicators provide the first stage of evaluation, and comprise of a number of independently measured values. They help to improve the implementation of the resource allocation mechanisms.

For further evaluation, it is obligatory to take into account a set of characteristics that are not directly comparable because these characteristics correspond to variables of different dimensions and unit of measurement. Thus, they have to be made comparable, e.g. by normalization, and then grouped into indicators. An indicator is defined as a ratio (a value on a scale of measurement) derived from a series of observed facts, which can reveal relative changes as a function of time. They provide feedback on system behaviour and also allow the analysis of performance and predictions of future performance.

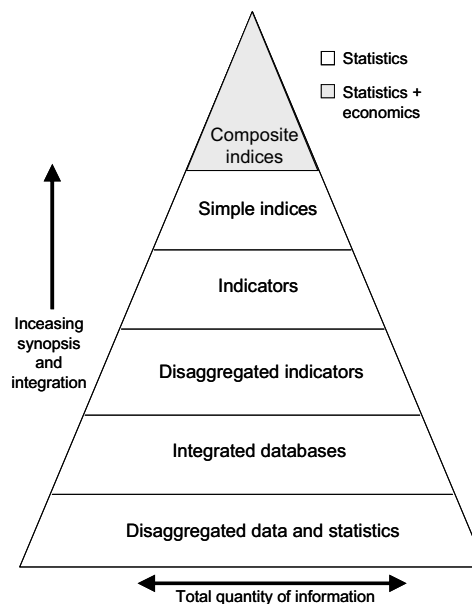


Figure 6: General method to obtain a composite index

Finally, the simple and composite indices are computed, which represent performance benchmarks. They express information in ways that are directly relevant to the decision-making process. In general, indicators help the assessment, the evaluation, and most important, they help to improve accountability. Although highly aggregated indices are attractive because of their simplicity, they also carry risks. Most of all, aggregates tend to mask real-life complexity and detail for policy-making. Highly aggregated indices are important in giving a view of overall progress, but they should be readily disaggregated into its components that may help to specify reasons for the index going up or down and also answer questions of interest to decision-makers working on lower scales. The pyramid with its different layers supports this analysis progress, it enables both analysis at lower scales and decision-making on highly aggregated indices.

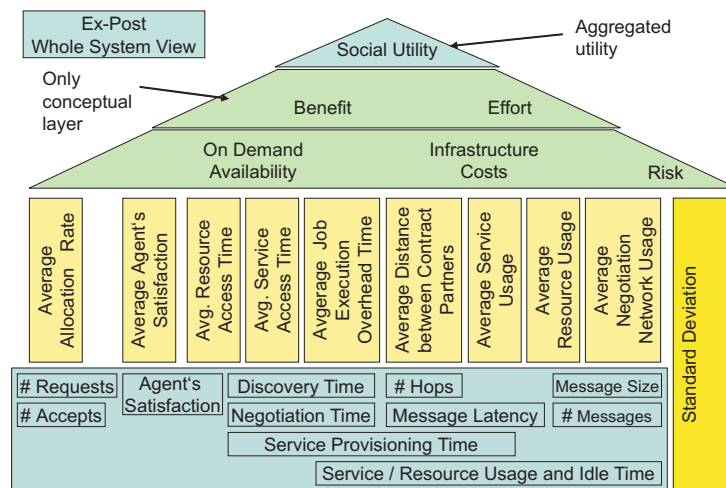


Figure 7: Metrics Pyramid

To support both technical and economic parameters, the evaluation process is divided into two layers. The first uses basic statistical concepts, while the second uses economic principles. The technical parameters, at the lower end of this pyramid, provide the basis for economic parameters that lie higher up in the pyramid, as illustrated in Figure 7. A prerequisite is, that the key parameters at the bottom layer can be technically measured in the system to be analysed.

Aggregated indices may be generated out of several of these low level parameters. The parameters at the higher layers of the pyramid have economic semantics and therefore a higher significance. At the top of the pyramid there is a composite index defined as (global) “social utility” – providing a single metric to specify the overall performance of the allocation strategy being used in a particular application.

The selection of the technical parameters, and their aggregation to simple and composite indicators, can vary, depending on the resource allocation method and/or the measurable technical metrics of the application domain. The framework gives the possibility to add and remove metrics to optimize the framework for a special application domain without changing the general concepts used to build up this pyramid.

2.4.2 Technical and economical metrics

This section defines and illustrates various layers and aggregation steps of the metrics pyramid, as used in the project CATNETS to analyse centralized and decentralized economic resource allocation methods.

The approach to technical metrics focuses on providing generic, easy to measure parameters, which can subsequently be aggregated. Technical layer metrics can be classified into: (I) efficiency measures (e.g. number of requests, number of acceptances); (II) utility measures (e.g. agent satisfaction); (III) time metrics (e.g. discovery time, negotiation time, service provisioning time) which are measures of the rate of change of market processes; and message-based metrics (IV) to measure the activity of users to communicate to find resources and services. The technical metrics include:

- Number of Requests: This metric measures demand, counted as the number of requests for services and resources.
- Number of Acceptance: The number of acceptance measures successfully requested (acknowledged) services and resources.

- **Agent Satisfaction:** The individual agent satisfaction measures the utility gain of a single transaction, which is the distance between his lowest price willing to pay in this transaction and the final price of the agreement. It is defined as the ratio between the subjective reservation value and transaction price.
- **Discovery Time:** This metric is used to measure the time to find a set of possible negotiation partners.
- **Negotiation Time:** The negotiation time measures the time needed to finish the negotiation between one buyer and one or several sellers. The measurement of the negotiation time starts after service discovery has completed, and ends before service provisioning.
- **Service Provisioning Time (Effective Job Execution Time):** The evaluation framework defines the service provisioning time as the service usage time of one transaction. It optionally includes also setup times, etc.
- **Hops:** The number of hops describes the distance between a service consumer and a service provider, counting the number of hops a message needs to traverse between them (this may be averaged over all the messages exchanged between a consumer and a provider).
- **Message Latency:** The message latency measures the time a message needs to arrive at the communication partner. Message latency is a parameter that indicates the performance of the physical network link and its message routing capabilities. It is expected that a large distance (cf. also number of hops) between a service consumer and service provider should lead to higher message latency. Whether hops or latency is used to measure distance should be decided casually.
- **Message Size:** This metric is used to measure the message size (e.g. in kilobytes).
- **Number of Messages:** This value counts the number of messages needed for service allocation. The traffic on a physical network link could be computed by multiplying the message size with the number of messages sent.

These metrics are used in two aggregation steps.

Simple indicators (first aggregation)

The simple indicator layer defines a set of independent metrics which are normalized between 0.0 and 1.0. This makes it easier to find valid functions for the layers above, such as on demand availability and infrastructure cost. The technical metrics may be combined to obtain a framework that enables evaluation of different service oriented architectures. The aggregated metrics at the simple indicator layer are:

- **Allocation Rate:** This metric is a measure of the efficiency of the allocation process, which is computed using the number of requests and number of accepts. A buyer can demand services, but there is no guarantee that the allocation mechanism (centralized or decentralized exchange) performs a match between demand and supply.

$$allocation_rate = accepts/requests$$
- **Agent Satisfaction:** This metric implicitly shows the average surplus of the service provider or user (agent) in the system (normalized to the interval 0.0 and 1.0). A low value means that an agent has not been able to complete its goals successfully during the negotiation process. A high value means that the agent can constitute good results satisfying its requirements.

- **Access Time:** This indicator evaluates the time needed from the starting point of discovery until the final delivery of the service.
 $access_time = discovery_time + negotiation_time + provisioning_time$
- **Job Execution Overhead:** This is the additional time needed for negotiation. It refers to the overhead introduced during the service negotiation process. The overhead is the sum of the service access time and the resource access time.
 $job_execution_overhead = access_time_service + access_time_resource$
- **Distance between Contract Partners:** Message latency is the messaging time incurred by agents, and it is proportional to the distance between the sending and receiving nodes. The latter would be the ratio between the actual distance and the maximum distance between agents.
- **Service/Resource Usage:** The network usage will be evaluated by the ratio between provisioning time and the total simulation time. This evaluation would be conducted for both services and resources.
 $service_usage = provisioning_time / simulation_time$
 $resource_usage = provisioning_time / simulation_time$
- **Network Usage:** This metric is used to measure the total number of messages exchanged between two agents.

Composite indicators (second aggregation)

Composite indicators are an aggregation of simple indicators, and are closer in scope to the application than to the infrastructure. Simple indicators are normalized between a [0; 1] interval in order to be comparable with each other. From the economics perspective, composite indicators measure the benefit from exchange (on-demand availability), and costs incurred from activities in the market (infrastructure costs). An optimal allocation mechanism would have an on-demand availability value = 1, and infrastructure costs = 0. On Demand availability (ODM) is a composite indicator obtained as mean of simple indicators, and may be derived as:

$$ODM = 1/3 (allocation_rate + agent_satisfaction + job_execution_overhead)$$

Infrastructure Cost (IC) is calculated in the same way. It is also a mean of multiple simple indicators, and may be derived as:

$$IC = 1/4 (distance + service_usage + resource_usage + network_usage)$$

It may be possible to model some of these metrics as stochastic variables, giving a mean and standard deviation over which the given metric varies. In economic applications, variance would be a measure for the overall “risk” to achieve stability of a given metric development.

The social utility – A composite index

To be able to use economic concepts into the evaluation process, a virtual policy maker is introduced, who aims to optimize the allocation. The behaviour of the policy maker is described with a goal function which is to be minimized under some constraints. The constraints are determined by the structure of the economy, i.e. the underlying laws of the economy. The policy maker does not have complete knowledge of the economy, but uses a model of the economy based on statistical cause-and-effect relations. The policy maker collects all the metrics from the previous layer, has some preferences about ODM and IC and is aware about the distribution of profits and costs.

In order to follow this approach for ALNs, a loss function is built using the composite indicators ODM and IC. ODM and IC are taken as stochastic variables and their distribution across a population (the set of service users and providers). For the final evaluation of the scenarios, the first and second moments of ODM and IC have to be evaluated, whereas the first moment is the mean of ODM and IC, and second moment the variance of ODM and IC. In fact, a policy maker has preference to minimize the inverse of ODM, thereby leading to the minimization of costs and the minimization of variability, i.e. the fair distribution of social utility between agents. Applying the depicted metrics pyramid, a layered evaluation of different resource allocation methods is possible.

2.4.3 Statements on applying the Catallaxy to Application Layer Networks

The final statement “Catallaxy is more / less efficient than central mechanisms” could not be obtained from the experiments in general. However, there is a catallactic strategy configuration which achieves equal or better performance in terms of equal or better social utility values than the centralized allocation approach of the simulated scenarios. The following insights have been accomplished:

- The Catallactic mechanism needs to be improved in terms of messages needed. The total number of messages could be reduced by increasing the price step value in the current strategy implementation.
- The configuration complexity in terms of the number of parameters compared with other economic mechanisms is very high in the Catallactic mechanism. On the one hand, this flexibility enables to find a good parameter configuration in all evaluated scenarios. On the other hand, there is no single configuration which can be applied in most scenarios.
- Comparison with centralized mechanisms is difficult since the performance does depend on each catallactic agents` strategy in the catallactic scenarios. In the centralized case, the auctioneer`s decisions only depend on the incoming supply / demand messages. Catallactic agents follow a heuristic strategy whereas the centralized auctioneer implements a mathematical algorithm with theoretical foundations (state-of-the-art matching mechanism).
- Advantages of the centralized mechanism originate in the significantly lower number of messages needed and the fact that the catallactic case supports only single-attribute negotiations can take place, whereas the auctioneer mechanism can handle multi-attribute negotiations. This limitation was imposed by the agents` strategy implementation. The learning and decision algorithms were implemented to cover only single-attribute negotiations. For this project a model was used in which each agent could only be involved in one negotiation at a time. This model posed serious limitations on performance – for a more efficient system the option of parallel negotiations would have to be implemented. This model also results in a very restrictive blocking policy of the agents.

- Implementation aspects of the simulator were underestimated. More time was spent on testing of the simulator tools and the implemented agents. Therefore, the co-allocation implementation and the shared resource model have still experimental status.
- Analyzing the trading agents in the decentralized strategy, two dominant trading strategies lead to profit (positive fitness) in a bilateral negotiation. Either the agents make often concessions with small steps or the agents follow a strategy with a low concession rate together with a high step size. Weighting the last agreements with 60% seems a good value for the successful agents.
- No clear picture evidences the applied service and resource distributions in scenarios with the catallactic strategy. A placement of resource on good connected nodes results in better social utility values. This is an argument for organizing computing centers in a centralized way. Uniformly distributed service and resources lead to good availability of the specific service and resource. But, higher deviations from the mean values could be observed than in other experiments. The competition for several service and resource increases the social utility index and the system's loss.
- In absolute numbers, the decentralized allocation approach is able to achieve a social utility index value of around 0.4. This value is 20 % better than the measured value for the centralized allocation approach in the comparison scenario. Again, this indicates the strong influence of the configuration on the allocation performance of the catallactic approach.
- In the evaluated scenarios, the decentralized allocation approach shows higher allocation rates than the centralized allocation approach. The centralized method sacrifices a higher allocation rate for the economic outcome of the auction.
- Varying discovery timeouts increase the availability of the sellers and point at successful refinement of the decentralized search strategy.
- The implemented decentralized bargaining protocol is error-prone which exhibits the failure experiments. At the time of development, the main focus was more at a sound and functional messaging structure and implementation than a failure-resistant bargaining protocol. More effort is needed for improving the bargaining protocol.
- The metric pyramid and the selected metrics enable the analysis of complex application layer networks. Changes in the measured metrics are mapped very well to the upper layers of the pyramid. The aggregation steps keep the characteristics of the raw data and indicate different allocation performances in the evaluated scenarios. The general applicability of the metrics pyramid offers a broad evaluation of resource allocation approaches in future application layer networks. However, the metrics pyramid exposes limitations. The number of observations and the agent population should be evaluated in parallel to verify the expressiveness of the aggregated indexes.

- The decentralized allocation performance does not decrease in increasing network sizes. The hop count and discovery timeout parameters control the accessible area of the network for the decentralized allocation approach. Only a subset of (theoretically) available service and resources can be selected as trading partners. This helps achieving good social utility index values in large networks.
- Varying bandwidth decreases the system loss for the decentralized allocation algorithm. The bandwidth controls the set of available sellers. Not reachable seller can provide their service to another buyer which is located closer the requestor. Also, the set of unique trading agents increase. A more distributed behavior is observed. However, too low bandwidth increases the system's loss significantly.
- The hop count parameter which can be seen as a time-to-live parameter in P2P networks shows significant influence on the decentralized allocation approach. The following rule is derived from the experiments: The higher the service density is, the lower the hop count parameter should be selected.
- Randomizing the initial genotype values lead to lower performance of the decentralized allocation approach because the agents need time to adapt to the dominant strategies described above. Compared to simulations runs with one predefined genotype for all agents, a significant lower performance is measured.
- A crucial part of the catallactic strategy is the adaptation of the negotiation intervals for the the next negotiation which depends on the market price estimation. The selected dynamic strategy exhibits good performance in all simulation scenarios. It is expected that an improved adaptation strategy can lead to better performance in terms of a decreasing message number and better allocation times.

2.5 Future Research on Properties of Catallaxy applied to Computer Networks

From a computational and engineering perspective, Hayek's Catallaxy is successfully applied to computer networks in the CATNETS project. Explicit specifications of the Catallaxy for technical systems should encompass these agent features:

- **Adaptivity:** Agents learn from their own experience and from previous agreements, there is genetic recombination, mutation and selection, and agents are reactive and opportunistic being able to adapt their goals to local unpredictable and evolving environments. The current catallactic strategy implemented the first two adaptivity topics. The third issue gives further research prospects. More parameters could be taken into account for dynamic adaptation by the agent itself instead of being predefined at experiment start.
- **Autonomy and initiative:** The complex service handles the task to be executed. If it finds an opportunity (a seller which fits the requirements) – proactively

without the direct control of the user application – the agent delegates not only a specific task but also an objective to bring out in any way; the agent will find its way on the basis of its own learning and adaptation, its own local knowledge, its own competence and reasoning, problem solving and discretion. The results of the negotiation analysis demonstrate the successful application of the catallactic strategy to the CATNETS scenario. Agents gain profit depending on their own learning and adaptation capabilities, and their local knowledge. Of course, not all agents are successful in making profit. They don't have the knowledge for finding satisfying goals. Dissatisfaction is related to the agents' goals, which Hayek calls "the conscious purpose". The agents' level of wealth, knowledge and consumption determines the "conscious purpose" as a function of which the agent will be more or less satisfied, and more or less responsive to existing circumstances. If the agent knows only other agents with low fitness and is unaware of what else is possible, the agent will act within its local environment and isn't able to escape from it. There could be agents in the application layer network, which never gain profit. Further research should address finding a lower bound for the agent success for the catallactic strategy.

- Distribution and decentralization: In CATNETS, the multi-agent systems are open and decentralized. As the results of the simulations runs evidence, it is neither established nor predictable which agents will be involved in trades. There are no regional analysis tools of the agents and its trading partners available. Placing resource on good connected and varying bandwidth on the links increases the distribution and decentralization of the multi-agent system. Further experiments are needed to find out more rules for increasing the agents' distribution and decentralization. The execution times of the tasks are assumed to be constant in CATNETS. Agents may remain open during execution by means of being reactive to incoming inputs and to the dynamics of its internal state (for example, resource shortage or preference change). Further research on the Catallaxy could take this into account increasing the reactivity of the agents. The workflow simulations with a complex service requesting a sequence of basic service give prospects for further analysis, because nobody in the system knows the complete plan. Even the complex services don't know to which basic service a task on second or third place could be delegated. Policies are needed to ensure that the task of long workflows could still be executed because longer workflows come along with an increased failure probability.

The products traded on the service market are homogeneous products which are completely standardized. The experiment runs of the lessons learnt from CATNETS show, the implemented centralized and decentralized allocation approaches handle these products very well. The implemented Hayek's Catallaxy can even outperform the centralized auctions. The introduction of heterogeneous products enables a more realistic service picture. Also the users or applications of the catallactic market profit from heterogeneous services. They can specify their workflows with more details like service quality levels, priorities, etc. This leads to a broader applicability of the Catallactic strategy in new application domains and real world user requests.

As mentioned above, CATNETS assumes an open application layer market implemented by an open multi-agent system. The traded quality of the products in terms of execution time is always constant. No service quality changes are taken into account.

Further research prospects could extend the CATNETS model with changing service quality which is required by future interactive Grid and P2P application based on service-oriented architecture. Additionally, shared resource models have to guarantee certain quality levels to users. Hayek's concept of a "spontaneous order" could be eroded by cheating agents in such environments. There is no institution which forces the agents to be honest. The Catalaxy could benefit by the introduction of electronic institution and concepts from social control like reputation mechanisms. This comes along with a change from single attributive decision making to multi-attributive decision making and reasoning. The agents have to decide how much risk they agree to and how they should adapt their preference structure in case of risky services and resources. Besides using social concepts, these risky assets can also be handled by economic risk management. Concepts from risk transfer like insurances or risk mitigation like portfolio optimization could help to reduce the financial risk of the agents. It would be interesting to see how the influence of the social and economic approaches influences the CATNETS metrics pyramid for evaluation of Hayek's Catalaxy.

Future hardware developments will soon make possible the construction of very large scale (one million of agents and above) models that obviate the need for representative agents. Artificial agent communities and economies of such scale need more research to overcome the lack of understanding of realistic behaviour of agents and institutions. Hayek's Catalaxy is one important concept to help understanding such complex and dynamic agent communities. We identified some future research prospects in the field of artificial economies:

- The markets could benefit from the introduction of new agent roles. Intermediate traders which buy and sell basic services increase the liquidity of the markets. More trading on the markets will lead to more exchanged price signals. The agents' learning and reasoning components could gain better price estimations.
- The price signals which are exchanged on the markets should be honest and unaltered. An independent monitoring institution could ensure trusting price signals. Traders pay for this market monitoring service. A loss of confidence would force the traders to select another trading partner.

2.6 Conclusions

In Application Layer Networks like future Grid and P2P networks, optimal resource allocation has to be carried out in two dimensions. One is the maximization of the utilization of technical resources. The aim in this dimension is - independently of the economic incentive structure - to carry out a load balancing on heterogeneous and distributed computational resources. This guarantees that all resources are used and no resources are "wasted" while being idle. In contrary to this allocation paradigm, the economic resource allocation is aligning the deployment of resources along the economic utility of the individual Grid nodes. Mechanisms like multi-attribute combinatorial exchanges enable an incentive compatible, efficient, individual rational and computational tractable way of allocating these resources. This may imply that resources stay idle in case the willingness to pay of the resource requester does not reach the reservation price stated by the resource owner. This effect is demonstrated by the experiment runs: The centralized resource allocation approach shows lower allocation rates than the catalactic allocation method. However, it enables an economically sound construction of allocation for Grid resources.

Within the CATNETS project, two dimensions of such markets have been investigated: One is the application of decentralized - Catallaxy-based - market mechanisms, which uses flooding for resource discovery and an iterative bilateral bargaining protocol for negotiation of resources. The alternative is the application of classic institutional - or centralized central - market mechanisms. If properly designed, markets implement an economical efficient allocation of resources. Such multi-attribute multi-unit mechanisms to ALN/Grid allocations have been developed as a benchmark. They provide an allocation for Grid resources up to a certain size. Mechanisms are developed that enable a two-tiered allocation of Grid resources, which fulfill most of the above-mentioned desiderata. In the first tier of these markets, service consumer, who may not have a clue of what kind of computational resources they need, can trade with service providers about their service needs. These service providers themselves then act on a second market tier - the resource market - where they purchase the resources they need in order to carry out the services. These two markets are interrelated through the price that is determined in the first tier. State of the art in this research field is that incentive compatible, allocative efficient and individual rational allocation mechanisms are identified. However, these mechanisms are very complex to compute for large number of market participants (they are NP-complete) and hence not applicable on large-scale setups while the decentralized market mechanism still achieves performant results in large settings.

The comparison of the two approaches shows that:

- The centralized mechanism used a significantly smaller number of messages for communication. The bargaining of the catallactic mechanism had a significantly higher cost in terms of messages exchanged. The number of messages in the implemented auctions remains always the same for one allocation, whereas the number of messages varies along the iterative negotiation steps in the decentralized allocation approach. By limiting the number of negotiation rounds, a lower number of messages can be achieved.
- In the catallactic mechanism, only single-attribute negotiations could take place, whereas the auctioneer mechanism of the centralized approach could handle multi-attribute negotiations. In our particular case, this limitation was imposed by the agents` strategy implementation. The learning and decision algorithms were implemented to cover only single-attribute negotiations.
- The implementation of the agents was based on a model in which each Catallactic agent could only be involved in one negotiation at a time. This model poses serious limitations on performance. A more efficient model with parallel negotiations isn't implemented due to its additional complexity. The same limitation holds with the centralized allocation approach. Parallel bidding is not supported.
- The potential of Catallaxy as resource allocation mechanism has been motivated by observing its usage in many real life situations. From the project results, we have noted that the complex models which are behind this usage, however, have not shown to be easily portable into artificial systems.
- The use of Catallaxy in daily life situations requires very sophisticated capabilities, which are not easily reproducible by technical implementations. Providing

each artificial agent with the "intelligence" needed to work in such a resource allocation approach has currently a high cost and needs an expert user.

- The contact with the system, the information dissemination and lookup, the interaction with other participants require very challenging technical solutions. It is considered difficult to provide technical solutions based on standard toolkits which could provide reliable functionality without limitations in large-scale scenarios and under realistic conditions.
- The evaluated scenarios have shown to be successfully realizable with less complex solutions than Catallaxy like a centralized approach based on auctions. The limitations of the centralized solution have not been identified in the scope of the requirements. Potential bottlenecks like the centralized components appear to have easier solutions for overcoming them than applying Catallaxy.
- We cannot exclude the possibility that particular scenarios which might arise from future distributed applications could provide circumstances in which Catallaxy remains an interesting option for providing a solution. Providing the building blocks of Catallaxy in a useable way, achieve a feasible configuration by the users, however, would ease any practical application of the approach.
- A final statement on the performance of the Catallaxy approach compared to centralized auction approach is difficult to obtain. A main problem lies in the technical aspects of the implementation. Catallaxy works best in large scale scenarios, but a sufficient simulation for Catallaxy needs larger technical resources. On the other hand, the simulation time needed for the centralized approach increases dramatically with growing simulation size.
- From an implementation perspective, even the moderate complexity of the heuristic bargaining strategy leads to a noticeable variance of the simulation results, when compared with the predictable results of the auctioneer's algorithm. The calibration of the simulation and a working prototype became an important task in the CATNETS project.
- With regard to the Grid market parameters, we have achieved various possibilities for adaptation to real world settings. Virtualized resources and resource bundles are supported.
- The final statement "Catallaxy is more / less efficient than central mechanisms" could not be obtained from the experiments in general. However, there is a catallactic strategy configuration which achieves equal or better performance in terms of equal or lower social utility values (that means better score for social welfare) than the centralized allocation approach of the simulated scenarios.

The recommendation for future research in this domain is threefold:

1. More research effort should be devoted to the combination between the communities that do research in the technical allocation and conceptualization of application layer networks as well as the economic allocation in future Grid

and P2P networks. Both sides can - while commercializing the application layer networks - not continue without the other.

2. Besides technical standardization, efforts for the development of economic standards and interaction schemes, e.g. based on Web Services or other concepts from Service Oriented Architectures, should be fostered for the practical utilization of future e-Infrastructures and e-Science.
3. Efforts for starting real-life pilots for Grid/P2P business models and Grid/P2P markets should be fostered, where researchers from computer science, economics and business administration commonly work on dynamic, economically sound and vertically integrated business concepts for the dynamic utilization of application layer networks and other e-Infrastructures.

3 Publishable Exploitable Knowledge and its Use

The expertise of project partners contributes links to key external research groups. Cardiff University has a joint project on "Grid Service Brokers" with Monash University (Australia), which is funded by the Victorian Partnership for Advanced Computing (VPAC) in Melbourne (Australia) and provides a link for collaboration between Rajkumar Buyya (Monash University) and Omer Rana (UWC).

ITC-irst participated in the DATAGRID project, which brings to the project a link with a large European real-world advanced grid and the possibility of direct access to additional large test-beds.

Universitat Politecnica de Catalunya has links with the CEPBA (European Center for the Parallelization of Barcelona) a research, development and innovation center with expertise in highperformance computing and infrastructure for grid applications with a joint research institute with IBM (CIRI). The group at UPC is coordinator of a 3 year (2002-2005) Spanish National integrated research project on Grid and Peer-to-Peer Middleware for Cooperative Learning with three universities and four Spanish SME involved.

Universitat Politecnica de Catalunya participates in the SORMA project. The catallactic middleware architecture of the CATNETS project is refined for the SORMA project to enable a mature open Grid market platform.

University of Karlsruhe coordinates the FP6-IST-GRID project SORMA, beginning from September 2006. The topic of this project is the construction of a Grid market platform, the evaluation of particular functionalities within the platform (such as billing & accounting), and the evaluation of business models for Grid environments.

University of Karlsruhe joined the German D-Grid initiative with the project "Biz2Grid: Moving Business to the Grid – An Application for the Automotive Industry". The centralized allocation approach developed in the CATNETS project will be used in the Biz2Grid project and extended by billing and account components. Biz2Grid offers a price based scheduler to the D-Grid community and to other related project. The virtues of the D-Grid extensions will be demonstrated by means of two commercial scenarios within the automotive industry. (<http://www.biz2grid.de>)

University of Bayreuth participates as an external consultant in the “The Utility Grid Project: Autonomic and Utility-oriented Global Grids for Powering Emerging E-Research Applications”. This project is funded from 2006 – 2009 by the Australian government. One milestone of this project is the implementation of a Catallaxy demonstrator. GridSim was extended to support the CATNETS scenario of two interconnected markets. The next release of GridSim will contain the catallactic strategy implementation developed by the CATNETS project.

University of Bayreuth is a participant of the EU project eRep and uses the CATNETS simulator in this project. The formal concept of electronic institutions and the notion of reputations will be introduced to the CATNETS simulator. Additionally, the support of BDI agents will be added to the CATNETS simulator which enables a richer support for more complex agent behaviour.

University of Bayreuth participates in the EU project SORMA. Empirical studies on the commercial application of Grid markets for SME will be accomplished. This helps to identify new application domains for commercial Grid markets.

University of Cardiff starts collaboration with the SURA/SCOOP project in the US (<http://scoop.sura.org/>). This project is investigating the use of pre-annotated data services being made available at different US locations, and using these in an aggregated manner to undertake data analysis to forecast possible hazards to coastal communities in the US. The project is already evaluating the use of the components being made available in the workflow engine as part of the CATNETS project, and further investigation on how pricing mechanisms can be used as a basis for selecting between different data sets is being investigated. Price in this context is derived using Quality of Service attributes. This work is being undertaken with Professor Gabrielle Allen at Louisiana State University.

University of Cardiff starts collaboration with UK VizNet to better understand how a market for visualization services can be supported for semi-real time data analysis and computational steering. This work is being undertaken with Professor Nick Avis at Cardiff University.

University of Cardiff:

The development of the prototype applications and use of WS-Agreement has provided a useful test case for the Grid application community interested in utilizing Service Level Agreements. This dissemination activity is being undertaken alongside participants from the GRAAP working group at the Open Grid Forum, and alongside the EU Technical Working Groups in Grid computing area.

University of Cardiff:

The applications developed also provide useful case studies for the use of a decentralized, market-based mechanism for service discovery in service-oriented architectures. Material from this work will be utilized in conference tutorials and lecture courses at Cardiff (such as the Distributed Systems course delivered to final year undergraduate (course code CM0355) and MSc students (course code CMT590)), and specifically the tutorial on “Autonomic Grid Computing” at the IEEE eScience conference in Bangalore (India) in December 2007.

4 Dissemination and Use

4.1 Dissemination to the Scientific Community

4.1.1 Organization of Scientific Meetings

Members of the CATNETS consortium organize workshops at conferences of the GRID computing and distributed artificial intelligence (DAI) research community to connect these different research fields, and in the intersection between DAI and Economics (agent-based computational economics). These workshops address new research aspects of agent technology in GRID environments. For further information, please see the information given on the websites of the scientific meetings. The organized events are:

- CCGrid 2005 Conference, Cardiff (Omer Rana (UWC) et al.: organizer)
- AGE (Agent-based Grid Economics) Workshop at CCGrid 2005 (<http://www.iw.uni-karlsruhe.de/ccgrid/>) (Daniel Veit (UKA), Björn Schnizler (UKA), Torsten Eymann (UBT): organisers)
- SGT (Smart Grid Technologies) Workshop at AAMAS 2005 (<http://www.aamas2005.nl/workshops.php>) (Daniel Veit (UKA), Torsten Eymann (UBT), et al.: organisers)
- SAACS (International Workshop on Self-Adaptive and Autonomic Computing Systems) in 2004 (Zaragoza) and 2005 (Copenhagen) (<http://cms1.gre.ac.uk/conferences/dexa2004/ws-2004/etc/themes-topics-saacs.htm>) (Torsten Eymann (UBT): member of steering committee, publications chair)
- EXYSTENCE Complex Systems Network of Excellence Thematic Institute, in Ancona (<http://www.complexityscience.org/index.php>) (Mauro Gallegati (UPM): organiser)
- Workshop on Agent Based Grid Computing, CCGrid 2006, Singapur (<http://users.cs.cf.ac.uk/O.F.Rana/agc2006/>) (Liviu Joita (UWC), Björn Schnizler (UKA), Pablo Chacín (UPC): organisers)
- Second International Workshop on Smart Grid Technologies (SGT-2006), ICAC 2006, Dublin (<http://www.caip.rutgers.edu/icac2006/>) (Torsten Eymann (UBT), Daniel Veit (UKA) et al: organisers)
- SAACS (4th International Workshop on Self-Adaptive and Autonomic Computing Systems), 2006 in Krakow, Poland. (Torsten Eymann (UBT): member of steering committee, publications chair)
- SOAS (2nd Conf. on Self-Organizing and Autonomic Systems, Sept 2006, Erfurt, Germany). Torsten Eymann, UBT: PC member, Omer F. Rana, UWC: invited speaker.
- International Workshop on Agent-based Grid Computing at 7th IEEE International Symposium on Cluster Computing and the Grid - CCGrid 2007, May 14-17, 2007, Rio de Janeiro, Brazil
- 8th International Conference on Business Information Systems (Wirtschaftsinformatik 2007), 28. 02. – 2. 03. 07, Karlsruhe, Germany
- Vice-Chair: 4th International Workshop on Grid Economics and Business Models (GECON 2007) at 13th Euro-Par Conference, IRISA/ENS Cachan, Rennes, France, August 28, 2007
- Program Committee Member: Workshop on Service Level Agreements, IEEE Grid 2007, September 19-21, Austin, Texas, USA

- Program Committee Member: The 4th International Conference on Autonomic and Trusted Computing (ATC-07), Hong Kong, China, July 11-13, 2007
- Program Committee Member: 2nd International Workshop on Engineering Emergence in Decentralised Autonomic Systems to be held at the 4th IEEE International Conference on Autonomic Computing (ICAC 2007) June 11-15, 2007, Jacksonville, Florida, USA
- Program Committee Member: Workshop on Complex Adaptive Economic Systems at the 9th European Conference on Artificial Life, September 10th-14th, 2007, Lisbon, Portugal
- Steering Committee Member: Workshop on Agent based Grid Computing at 7th IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2007), Rio de Janeiro, Brazil, 14-17 May 2007
- Invited Panel Disputant: IEEE International Conference on Pervasive Services 2006 (ICPS), Lyon, France, June 2006
- Program Committee Member: The 2006 International Conference on High Performance Computing and Communications, Special Session on Service Level Agreements, München, Deutschland, September 2006
- Program Committee Member: 3rd International IEEE Conference on Autonomic and Trusted Computing, Wuhan and Three Gorges, China, September 2006
- SGT (Smart Grid Technologies) Workshop at AAMAS 2008 (Daniel Veit (UKA), Torsten Eymann (UBT), et al.: organisers)

4.1.2 Refereed Journal Publications

Major results of the CATNETS project will be published in scientific journals. For the time being, we have the following contributions accepted for publication:

1. Torsten Eymann, Michael Reinicke, Felix Freitag, Leandro Navarro, Oscar Ardaiz, Pau Artigas: "A Hayekian Self-Organization Approach to Service Allocation in Computing Systems", *Journal of Advanced Engineering Informatics*, Elsevier, 2005.
2. Reinicke, M.; Streitberger, W.; Eymann, T.: "Evaluation of Service Selection Procedures in Service Oriented Computing Networks". In: *Multi Agents and Grid Systems*, Vol. 1, 2005.
3. Ardaiz, O.; Catalano, M.; Chacin, P.; Chao, I. ; Eymann, T.; Freitag, F.; Gallegati, M.; Giulioni, G.; Joita, L.; Navarro, L.; Neumann, D.; Rana, O.; Reinicke, M.; Schiaffino, S. B., Ruben Carvajal; Streitberger, W.; Veit, D.; Zini, F.: "Catallaxy-based Grid Markets". In: *Multi Agents and Grid Systems*, 2005.
4. Eymann, T., Reinicke, M.; Freitag, F.; Navarro, L.; Ardaiz, O.; Artigas, P.: A Hayekian Self-Organization Approach to Resource Allocation in Autonomic Computing Systems. In: *Journal of Advanced Engineering Informatics*, Vol. 20, 2006.
5. S. Ludwig, O. Rana, D. Veit (2005) MatchMaking Shell: A Shell for supporting Service Negotiation In: *Journal of Decision Systems (JDS) 13/2004*. *Electronic Negotiations*, p. 461 - 475.
6. D.G. Cameron, R. Carvajal-Schiaffino, C. Nicholson, K. Stockinger, F. Zini, A.P. Millar, and L. Serafini. Formal Analysis of an Agent-based Optimisation Strategy for Data Grids. *Multiagent and Grid Systems*, 2(2), 2006.
7. B. Schnizler, D. Neumann, D. Veit, C. Weinhardt (2006), Trading Grid Services - A Multi-attribute Combinatorial Approach, *European Journal of Operational Research*, forthcoming

8. D. Neumann, D. Veit, Ch. Weinhardt (2006), Grid Economics: Market Mechanisms for Grid Markets, In: Grid Computing: Konzepte, Technologien, Anwendungen, Th. Barth, A. Schüll (Hrsg.), Vieweg Verlag, pp. 64-83
9. D. Veit (2006) Markets for Non-Storable Goods - Analyses under Application of Agent-Based Simulations, Kumulative Habilitationsschrift an der Fakultät für Wirtschaftswissenschaften der Universität Karlsruhe (TH)
10. D. Veit, T. Eymann, N. Jennings, J. P. Müller (2005), Guest Editorial Special Issue on "Smart Grid Technologies & Market Models", In: Multiagent and Grid Systems – An International Journal 1(4) (2005) 235–236.
11. W. Streitberger, S. Hudert, T. Eymann, B. Schnizler, F. Zini, M. Catalano (2007), On the simulation of Grid market coordination approaches, Journal of Grid Computing, forthcoming
12. C. Weinhardt, B. Schnizler, S. Luckner (2007), Market Engineering, Group Decision and Negotiation 2007, Montreal, Canada
13. Chao I., Brunner R., Freitag F., Navarro L., Chacin P., Ardaiz O., Joita L., and Rana O.F. (2007) - Decentralized Grid Market Infrastructure for Service Oriented Grids, accepted in the 2nd review process to the Special Issue (01/2008) on Service-Oriented Architectures and Web Services, in Journal WIRTSCHAFTSINFORMATIK (Business Informatics), Germany, 2008
14. Joita L, Rana O.F., Chacin P., Chao I., Freitag F., Navarro L., Ardaiz A. (2007) - "A Catallactic Market for Data Mining Services", in International Journal of Future Generation Computer Systems (FGCS) - Grid Computing: Theory, Methods & Applications, Volume 23, Issue 1, January 2007, ISSN 0167-739x, pp. 146-153. www.sciencedirect.com
15. Joita L, Rana O.F., Chacin P., Chao I., Freitag F., Navarro L., Ardaiz A. (2006) - "Application Deployment on Catallactic Grid Middleware", in Journal of IEEE Distributed System Online, ISSN: 1541-4922, vol. 7, no. 12, 2006, art. no. 0612-oz001, p. 1-17, <http://dsonline.computer.org>

4.1.3 Publications & Attendance at Scientific Meetings

Due to the wide spectrum of domains covered by the project, there are a large number of publication opportunities. At this stage of the project the targeted events with publications and attendance of project members are:

Submitted and accepted (conferences):

1. Reinicke, M.; Streitberger, W.; Eymann, T.: "Evaluation of Service Selection Procedures in Service Oriented Computing Networks". In: Proceedings of the First International Workshop on Smart Grid Technologies (SGT05), Utrecht, Niederlande, 2005.
2. Eymann, T.; Ardaiz, O.; Catalano, M.; Chacin, P.; Chao, I.; Freitag, F.; Gallegati, M.; Giulioni, G.; Joita, L.; Navarro, L.; Neumann, D.; Rana, O.; Reinicke, M.; Schiaffino, S. B., Ruben Carvajal; Streitberger, W.; Veit, D.; Zini, F.: "Catallaxy-based Grid Markets". In: Proceedings of the First International Workshop on Smart Grid Technologies (SGT05), Utrecht, Niederlande, 2005.
3. M. Reinicke (2005) Eine ökonomische Bewertung der Dienstausswahlverfahren in serviceorientierten Overlaynetzwerken Wirtschaftsinformatik 2005, Bamberg, Germany

4. D. Neumann, B. Schnizler (2005) Transforming Auction Mechanisms into Protocols, Group Decision and Negotiation 2005, Vienna, Austria
5. M. Grunenberg, B. Schnizler, D. Veit, C. Weinhardt (2005) Innovative Handelssysteme für Finanzmärkte und das Computational Grid, 67. wissenschaftliche Jahrestagung des Verbandes der Hochschullehrer für Betriebswirtschaft, Kiel, 2005
6. B. Schnizler (2005) Market Based Coordination in the Computational Grid - A Market Engineering Approach, Beitrag für das Doktorandenkolloquium der Wirtschaftsinformatik, 2005
7. B. Schnizler, D. Neumann, C. Weinhardt (2004) Resource Allocation in Computational Grids - A Market Engineering Approach. Proceedings of the WeB 2004, Washington
8. D. Neumann, J. Maekioe, C. Weinhardt (2005) CAME - A Toolset for Configuring Electronic Markets. To appear in proceedings of the ECIS 2005, Regensburg
9. B. Schnizler, S. Luckner, C. Weinhardt (2006) , Automated Trading across E-Market Boundaries, S. Seifert and C. Weinhardt (Eds.): Group Decision and Negotiation (GDN) 2006
10. P. Chacin, L. Joita, B. Schnizler (2006), Flexible Architecture for Supporting Auctions in Grids, Proceedings of the Second International Workshop on Smart Grid Technologies (SGT) 2006
11. B. Schnizler, D. Neumann, D. Veit, C. Weinhardt (2006), Moving Markets to the Grid, 68. Wissenschaftliche Jahrestagung des Verbandes der Hochschullehrer für BWL
12. D. Neumann, S. Lamparter, B. Schnizler (2006), Automated Bidding for Trading Grid Services, Proceedings of the European Conference on Information Systems (ECIS), 2006
13. S. Lamparter, B. Schnizler (2006), Trading Services in Ontology-driven Markets, Proceedings of the 21st Annual ACM Symposium on Applied Computing, Dijon, France
14. T. Eymann, D. Neumann, M. Reinicke, B. Schnizler, W. Streitberger, D. Veit (2006) On the Design of a Two-Tiered Grid Market Structure, Business Applications of P2P and Grid Computing, MKWI 2006
15. Pablo Chacin, Liviu Joita, Björn Schnizler (2006) - "Flexible Architecture for Supporting Auctions in Grids", Proceedings of the 2nd International Workshop On Smart Grid Technologies 2006 (SGT2006) Workshop, In co-location with the 3rd International Conference on Autonomic Computer, 12-16 June 2006, Dublin, Ireland.
16. Liviu Joita, Omer F. Rana, Pablo Chacin, Isaac Chao, Felix Freitag, Leandro Navarro, Oscar Ardaiz (2005) - "Application Deployment using Catallactic Grid Middleware", Proceedings of the 3rd International Workshop on Middleware for Grid Computing (MGC 2005), co-located with ACM/USENIX/IFIP 6th International Middleware Conference 2005, 28 November - 2 December 2005, Grenoble, France. The paper has been nominated by the workshop organisers as one of the top two papers of the workshop.
17. Omer Rana, Jeremy Hilton, Liviu Joita, Peter Burnap, Jaspreet Singh Pahwa, John Miles, and W. Alex Gray (2005) - "Secure Virtual Organisations: Protocols and Requirements", at Information Security Solutions Europe (ISSE) 2005 - The Independent European ICT Security Conference and Exhibition, 27-29 September 2005, Budapest, Hungary; in ISSE 2005 - Securing Electronic Business Processes, by Sachar Paulus, Norbert Pohlmann, and Helmut Reimer, edited by Vieweg Verlag, 2005, ISBN 3-8348-0011-2, p. 422-432.

18. M. Reinicke, W. Streitberger, T. Eymann, M. Catalano, G. Giulioni (2006) Economic Evaluation Framework of Resource Allocation Methods in Service-Oriented Architectures, Proceedings of the 8th Conference on E-Commerce Technology (CEC06), San Francisco, June 26th-29th
19. C. Weinhardt, R. Studer, C. Holtmann, B. Schnizler, A. Ankolekar, N. Stojanovic (2006), Service Science, Management, Engineering and eOrganisations - A Position Paper, Proceedings of the SSME Summit: Education for the 21st Century, Palisades (USA)
20. B. Schnizler (2006), MACE: A Multi-Attribute Combinatorial Exchange, N. Jennings, G. Kersten, A. Ockenfels, C. Weinhardt (Eds.): Dagstuhl Seminar Proceedings 06461 Negotiation and Market Engineering, Internationales Begegnungs- und Forschungszentrum (IBFI), Schloss Dagstuhl, Germany
21. Eymann, T.; Streitberger, W.; Hudert, S.: "Global Grids - Making a Case for Self-Organization in large-scale Overlay Networks". In: Proceedings of the Second Trustworthy Global Computing Workshop (TGC 2006), Lucca, Italy, 2006
22. Streitberger, W.; Eymann, T.; Veit, D.; Catalano, M.; Giulioni, G.; Joita, L.; Rana, O. F.: "Evaluation of Economic Resource Allocation in Application Layer Networks – A Metrics Framework". In: 8. Internationale Tagung Wirtschaftsinformatik (WI 2007), Karlsruhe, Deutschland, 2007
23. B. Schnizler (2007), Resource Allocation in the Grid : A Market Engineering Approach, Universitätsverlag Karlsruhe, ISBN: 978-3-86644-165-1
24. A. Oberweis, C. Weinhardt, H. Gimpel, A. Koschmider, V. Pankratius, B. Schnizler (Eds.) (2007), eOrganisation: Service-, Prozess-, Market-Engineering, Vol. 1 & 2, Universitätsverlag Karlsruhe, ISBN: 978-3-86644-094-4 (Vol. 1), 978-3-86644-095-1 (Vol. 2), 978-3-86644-093-7 (Set)
25. Catalano, M; Eymann, T.; Giulioni, G.; Streitberger, W.: "Auctions vs. Bilateral Trading: New Assessment on Efficiency Performances". In: 13th International Conference on Computing in Economics & Finance (CEF 2007), Montreal, Canada, 2007.
26. D. Neumann, B. Schnizler, I. Weber, C. Weinhardt (2007), Second Best Combinatorial Auctions - The Case of the Pricing Per Column Mechanism, Proceedings of the 40th Annual Hawaii International Conference on System Sciences, Computer Society Press
27. Eymann, T.; Streitberger, W.; Hudert, S.: "CATNETS - Open Market Approaches for Self-Organizing Grid Resource Allocation". In: Proceedings of the 4th International Workshop on Grid Economics and Business Models (GECON 2007), Rennes, France, 2007.
28. Dias de Assuncao, M.; Streitberger, W.; Eymann, T.; Buyya, R.: "Enabling the Simulation of Service-Oriented Computing and Provisioning Policies for Autonomic Utility Grids". In: Proceedings of the 4th International Workshop on Grid Economics and Business Models (GECON 2007), Rennes, France, 2007.
29. Chao I., Ardaiz O., Sangüesa R., Joita L., and Rana O.F. (2007): Optimizing Decentralized Grid Markets through Group Selection, Workshop on Economic Models and Algorithms for Grid Systems, in conjunction with The 8th IEEE International Conference on Grid Computing (Grid 2007), Austin, Texas, USA, September 19 - 21, 2007.
30. Brunner R., Chao I., Chacin P., Freitag F., Navarro L., Ardaiz O., Joita L., and Rana O.F. (2007): Assessing a distributed market infrastructure for economics-based service selection, accepted to the International Conference on Grid computing, High-Performance and Distributed Applications (GADA'07), Vilamoura, Algarve, Portugal, November 29 - 30, 2007.

31. Felix Freitag, Pablo Chacin, Isaac Chao, Rene Brunner, Leandro Navarro, Oscar Ardaiz Performance Measuring Framework for Grid Market Middleware , accepted for 4th European Performance Engineering Workshop (EPEW 2007), Berlin, Germany, September 2007
32. Joita L., Rana O.F., Chao I., Chacin P., Freitag F., Navarro L., and Ardaiz O. (2007) - Service Level Agreements in Catallaxy-based Grid Markets, in Usage of Service Level Agreements in Grids Workshop, in conjunction with The 8th IEEE International Conference on Grid Computing (Grid 2007), Austin, Texas, USA, September 19 - 21, 2007.
33. J. Altmann, D. Veit (2007) Grid Economics and Business Models – Proceedings of the 4th International Workshop on Grid Economics and Business Models (GECON), Springer Lecture Notes in Computer Science (LNCS) 4685
34. B. Schnizler, D. Neumann, D. Veit, C. Weinhardt (2006) Trading Grid Services - A Multi-attribute Combinatorial Approach In: European Journal of Operational Research, in Press.

Attended Conferences and Workshops:

1. Wirtschaftsinformatik 2005, Bamberg
2. Workshop on eBusiness 2005, Washington
3. ICIS 2005, Washington
4. DEXA 2005, Zaragossa
5. Markets, Negotiations and Dispute Resolution in the New Economy 2005, Montreal
6. Doctoral Consortium Wirtschaftsinformatik, Bamberg
7. Global Grid Forum 11
8. "Computing and Markets" Seminar, Dagstuhl
(<http://www.dagstuhl.de/05011/Materials/>)
9. ECAI 2006, The 17th European Conference on Artificial Intelligence.
10. IEEE International Conference on Autonomic Computing, Dublin, Ireland
11. ICPS'06 : IEEE International Conference on Pervasive Services 2006
12. International Conference on Self-Organization and Adaptation of Multi-agent and Grid Systems (SOAS 2005), Glasgow, Schottland
13. GridKA School 2006, Karlsruhe, Germany
14. Group Decisions and Negotiations (GDN), 2006, Karlsruhe, Germany
15. Multikonferenz Wirtschaftsinformatik (MKWI) 2006, Passau, Germany
16. Market Engineering Day 2005, Karlsruhe Germany
17. 8th IEEE/ACM International Conference on Grid Computing (Grid 2007), Austin, Texas
18. ZUMA Simulation Workshop 2007, Koblenz, Germany
19. 4th International Workshop on Grid Economics and Business Models (GECON 2007), Rennes, France
20. CoreGrid Workshop 2007, Rennes, France
21. CoreGrid Summer School 2007, Budapest, Hungary
22. International Conference on Grid computing, High-Performance and Distributed Applications (GADA'07), Vilamoura, Algarve, Portugal
23. 13th International Conference on Computing in Economics & Finance (CEF 2007), Montreal, Canada
24. 8. Internationale Tagung Wirtschaftsinformatik (WI 2007), Karlsruhe, Deutschland
25. Dagstuhl Seminar „Negotiation and Market Engineering“, Internationales Begegnungs- und Forschungszentrum (IBFI), Schloss Dagstuhl, Germany

Invited Talks/Panel Participation:

1. (Daniel Veit) Invited Talk: Grid Business Models, eingeladener Vortrag bei der Eröffnung der GridKA School 2006, FZK, Karlsruhe, Germany, September 2006
2. (Daniel Veit) Panel Disputant: IEEE International Conference on Pervasive Services 2006 (ICPS), Lyon, France, June 2006
3. (Daniel Veit) Panel Disputant: International Conference on Self-Organization and Adaptation of Multi-agent and Grid Systems (SOAS 2005), Glasgow, Scotland, December 2005
4. (Werner Streitberger) Invited Talk: CATNETS Project, at HP Research Labs, Palo Alto, US, July 2006
5. (Omer Rana and Michael Reinicke) Invited talk: CATNETS Project, at GGF-16, Athens, Greece, February 2006
6. (Omer Rana) Invited talk: Managing Dynamic Service Level Agreements in CATNETs at GGF-17, Tokyo, Japan, May 2006
7. (Omer Rana) Invited talk: Center of Excellence Approach to collaborative working -- at UK-Singapore Collaboration Workshop, organised alongside GridAsia 2006, Singapore, May 2006
8. (Omer Rana) Invited talk: Grid Computing Workshop: Trends, Technologies and Collaborative Opportunities workshop, Bangkok, Thailand, May 2006
9. (Omer Rana) Invited talk: Creating and Managing Dynamic Service Level Agreements -- at NetObject.Days, Erfurt, Germany, September 2006
10. (Bjoern Schnizler) talk: MACE: A Multi-Attribute Combinatorial Exchange, Dagstuhl Seminars on Negotiation and Market Engineering 2006
11. (Werner Streitberger) talk: Simulation von Grid-Versicherungen. In: Doctorial Consortium der 8. Internationalen Tagung Wirtschaftsinformatik, Karlsruhe, Deutschland, 2007
12. (Daniel Veit) Invited Talk: Grid Economy and Business Models; from Web to Grids, Invited talk at the Open Workshop on E-Infrastructures of the E-IRG Reflection Group of the European Union within the German EU Presidency, 20.4.2007
13. (Daniel Veit) Invited Talk: Allocation of Resources in Service-Oriented Computing - Multi-attribute Combinatorial Trading, invited talk in the Department Seminar, Department of Industrial Engineering and Operations Research, University of California, Berkeley, USA, 5.02.2007
14. (Daniel Veit) Invited Talk: Computational Services as Utility - Multi-attribute Combinatorial Trading, invited talk at the Department of Management Science and Engineering der Stanford University, Palo Alto, USA, 1.02.2007
15. (Daniel Veit) Invited Talk: Grid Business Models, invited talk at the opening of the GridKA School 2006, Forschungszentrum Karlsruhe (FZK) Karlsruhe, Germany, 11.09.2006

4.1.4 Participation in Scientific Networks

The project participants are engaged also in other scientific networks, both nation- and EU-wide. Selections of those networks are:

AgentLink III was the premier Co-ordination Action for Agent Based Computing, funded by the European Commission's 6th Framework Program. Launched on 1st

January, 2004, it provided support for the network of European researchers and developers with a common interest in agent technology through events aimed at industry outreach, and standardisation issues, as well as providing support for academic events and providing resources through the AgentLink Portal. Participating in AL III were: UBT Bayreuth, UKA Karlsruhe, UPC Barcelona, Cardiff U.

The Complex Systems Network of Excellence, "Exystence", is funded by the European Commission to develop collaboration among European researchers interested in Complex Systems, from fundamental concepts to applications, and involving academia, business and industry. Participating in EXYSTENCE are: UBT Bayreuth, UPM Ancona.

PlanetLab is a global test bed for inventing and testing prototype Internet applications and services. The researchers aim to spark a new era of innovation by using "overlay" networks to upgrade and expand the Internet's features and capabilities. Being part of this scientific network gives the possibility to use PlanetLab for the prototype application. The results of CATNETS are expected to work towards improved resource allocation in PlanetLab and derived implementations. Participating in PlanetLab are: UPC Barcelona.

Agentcities (www.agentcities.org) is an initiative that aims create a next generation Internet that is based upon a worldwide network of interoperable services. These services, ranging from eCommerce to integrating business processes into a virtual organisation, can be accessed across the Internet, and have an explicit representation of the capabilities that they offer. Participating in AgentCities are: U Freiburg (prior partner, moved to UBT Bayreuth), UPC Barcelona.

AgentcitiesUK is an EPSRC-funded network: commencing in Autumn 2003 and completing in Autumn 2006. Participating in Agentcities.UK: Cardiff U.

DataGrid is a project funded by European Union. The objective is to build the next generation computing infrastructure providing intensive computation and analysis of shared large-scale databases, from hundreds of TeraBytes to PetaBytes, across widely distributed scientific communities. Participating in DataGrid is: ITC first Trento

The CATNETS project joined the CoreGrid network of excellence in order to further develop the ideas derived during project runtime together with a leading community or Grid Computing researchers. UBT participates in the resource allocation work package of CoreGrid.

4.2 Dissemination to a General Audience

The CATNETS consortium publishes its activities to a diverse audience. These activities are concentrated in the project website, project reports and newsletters, participation at exhibitions, presentations to industry and participation at standardization efforts.

4.2.1 The Website

The CATNETS website (<http://www.catnets.org>) was established in September 2004 (month 1 of the project). All information regarding the project is published on this website and is available in English language only.

The index page provides general information about the project CATNETS, the list of project partners and the contact address of the project manager.

A re-design process was started to extend the website and finished in June 2005. The new website delivers more information about the project. This process was done by the consortium member at the University of Karlsruhe.

The site is the main organ for public dissemination and the privileged support for information delivery on the project activities. The information provided on the web site encompasses:

- Presentation of the project (partners, objectives, planning).
- Presentation of public deliverables and main public results.
- Scientific presentation by the partners of the consortium.
- News about events like conferences, workshops, etc. related to the project.
- The source code and binary packages of the project including the simulator, the scenario generators, the CATNETS middleware, and the prototypes.

All project reports of the CATNETS project are public and are published on the CATNETS website. The CATNETS source code is released under the DATAGRID open source license which was updated for the CATNETS project.

4.3 Public deliverables

Deliverable 1.3: Annual Report of WP1 – Theoretical and Computational Basis
Daniel Veit, Georg Buss, Björn Schnizler, Dirk Neumann, Werner Streitberger, Torsten Eymann (2007)

Deliverable 2.3: Annual Report of WP2 – Simulation Framework
Werner Streitberger, Torsten Eymann, Floriano Zini, Björn Schnizler, Hong Tuan Kiet Vo (2007)

Deliverable 3.3: Annual Report on WP3: Proof-of-Concept Application
Oscar Ardaiz, Pablo Chacin, Isaac Chao, Felix Freitag, Liviu Joita, Leandro Navarro, Omer F. Rana, Werner Streitberger (2007)

Deliverable 4.3: Annual Report of WP4 - Performance Evaluation
Georg Buss, Michele Catalano, Pablo Chacin, Isaac Chao, Torsten Eymann, Felix Freitag, Liviu Joita, Leandro Navarro, Nils Parasie, Omer F. Rana, Björn Schnizler, Werner Streitberger, Daniel Veit (2007)

Deliverable 1.2: Annual Report of WP1
Daniel Veit, Georg Buss, Björn Schnizler, Dirk Neumann, Werner Streitberger, Torsten Eymann (2006)

Deliverable 2.2: Annual Report of WP2

Gaetano Calabrese, Bjoern Schnizler, Werner Streitberger, Floriano Zini (2006)

Deliverable 3.2: Annual Report on WP3: Proof-of-Concept Application

Oscar Ardaiz, Pablo Chacin, Isaac Chao, Juan Carlos Cruellas, Felix Freitag, Liviu Joita, Manuel Medina, Leandro Navarro, Omer F. Rana, Miguel Valero (2006)

Deliverable 4.2: Annual Report of WP4 - Performance Evaluation

Oscar Ardaiz, Michele Catalano, Pablo Chacin, Isaac Chao, Juan Carlos Cruellas, Felix Freitag, Liviu Joita, Manuel Medina, Leandro Navarro, Omer F. Rana, Bjoern Schnizler, Miguel Valero (2006)

Deliverable 1.1: Theoretical and Computational Basis

Björn Schnizler, Dirk Neumann, Daniel Veit, Michael Reinicke, Werner Streitberger, Torsten Eymann, Felix Freitag, Isaac Chao, Pablo Chacin (2005)

Deliverable 2.1: Analysis of Simulation Environment

Floriano Zini, Gianfranco Giulioni, Michael Reinicke, Werner Streitberger (2005)

Deliverable 3.1: Implementation of additional services for the economic enhanced platforms in Grid/P2P platform: Preparation of the concepts and mechanisms for implementation

Oscar Ardaiz, Pablo Chacin, Isaac Chao, Juan Carlos Cruellas, Felix Freitag, Liviu Joita, Manuel Medina, Leandro Navarro, Omer F. Rana, Miguel Valero (2005)

Deliverable 4.1: Evaluation and Metrics Framework

Michele Catalano, Gianfranco Giulioni, Werner Streitberger, Michael Reinicke, Torsten Eymann (2005)

Deliverable 0.1: Half Year Report of WP1

Björn Schnizler, Dirk Neumann, Daniel Veit, Björn Schnizler, Dirk Neumann, Daniel Veit, Michael Reinicke, Werner Streitberger, Torsten Eymann (2005)

Deliverable 0.2: Half Year Report of WP2

Torsten Eymann, Werner Streitberger, Michael Reinicke, Felix Freitag, Pablo Chacin, Isaac Chao, Björn Schnizler, Daniel Veit (2005)

Deliverable 0.3: Half Year Report of WP3

Felix Freitag, Pablo Chacin, Isaac Chao (2005)

Deliverable 0.4: Half Year Report of WP4

Gianfranco Giulioni, Floriano Zini (2005)

5 Co-ordinator Contact Details

Prof. Dr. Torsten Eymann
LS Wirtschaftsinformatik (BWL VII)
University of Bayreuth
95440 Bayreuth
Tel. +49 (921) 55-2807
Fax +49 (921) 55-2216
Email. catnets@uni-bayreuth.de
Website. <http://www.wi.uni-bayreuth.de>
CATNETS Website. <http://www.catnets.org>

This paper presents the final activity report of the CATNETS project. It summarizes the project objectives, the contractors involved, the work performed and the final results. A description of the dissemination activities concludes the report.