CATNET

IST-2001-34030

Evaluation of the Catallaxy paradigm for decentralized operation of dynamic application networks

Delivery 3: Catallaxy Evaluation Report

Covering period 2002-03-01 – 2003-02-28

Report Version: 0.27

Report Preparation Date: 2003-03-15

Classification: public

Contract Start Date: 2002-03-01 Duration: 12 months

Project Co-ordinator: Universitat Politècnica de Catalunya

Partners: Albert-Ludwigs-Universität Freiburg

Project funded by the European Community under the "Information Society Technologies" Programme (1998-2002)

Table of contents:

DELIVERABLE SUMMARY SHEET

Project Number: ITS-2001-34030

Project Acronym: CATNET

Title: Evaluation of the Catallaxy paradigm for decentralized operation of dynamic application networks

Deliverable N°: 3 – Catallaxy Evaluation Report

Due date: 03/2003

Delivery Date: 03/2003

Short Description:

The purpose of this deliverable D3 is to show, explain and discuss the results of the simulations conducted using the CATNET simulator as explained in deliverable D2 of this project. The experimental results of the CATNET project, especially with regard to whether the Catallaxy paradigm is able to provide sufficient resource allocation, will be compared using a hypotheses-based framework.

Chapter 1 explains the motivation, scenarios, criteria measured and the hypotheses to be investigated. Chapter 2 compares some basic experimental results by scenario and by measurement criteria. Chapter 3 shows how sensitive the results are to changes in the experimental setup. Chapter 4 summarizes and discusses the overall findings. Chapter 5 presents the achievements of this project and gives an outlook on possible future research directions.

Partners owning: Universitat Politècnica de Catalunya

Partners contributed: Albert-Ludwigs-Universität Freiburg

Made available to: Public

1. Introduction

The purpose of this deliverable D3 is to show, explain and discuss the results of the simulations conducted using the CATNET simulator as explained in deliverable D2 of this project. The experimental results of the CATNET project, especially with regard to whether the Catallaxy paradigm is able to provide sufficient resource allocation, will be compared using a hypotheses-based framework.

Chapter 1.1 outlines the motivation of the project. Chapter 1.2 explains the scenarios and the progress of the experiments. Chapter 1.3 explains the criteria measured by the simulator and shows their relevance with regard to realistic ALN scenarios. Chapter 1.4 details the hypotheses and sets the criteria in relation.

1.1 Motivation

Application-layer networks (ALN) are software architectures that allow the provisioning of services requiring a huge amount of resources by connecting large numbers of individual computers for information search, content download, parallel processing or data storage. Common concepts are Grid computing (mostly for distributed processing) and Peer-to-Peer-(P2P)-Computing (mostly for distributed data storage and access). In order to keep such a network operational, service control and resource allocation mechanisms are required. Here is an example to illustrate the type of networks:

Adobe's PDF file format is a common type for mixed text and graphics documents, mostly due to its preservation of layout specifics, in contrast to HTML. The files are created using the Acrobat Distiller service, which takes e.g. Microsoft Word or Postscript files and converts them. Usually, the Distiller service is installed locally and appears like an additional printer to the office application.
In addition, it is possible to submit a document to an Internet service like Adobe itself or T-Online. These services run on dedicated servers, convert the document and send a PDF file back. It is a centralized, classic client/server architecture approach.
However, in application-layer networks like the Grid or P2P networks, it is principally possible to provide the same service in completely decentralized fashion. Instead of one dedicated server instance, the relatively lightweight PDF conversion service could be provided redundantly by any network resource, and accessed "on-demand" by clients all over the network. The word-processor client programs would transparently address such a networked PDF conversion service instance in the background, without to disturb the user.
The main advantage for the normal user are costs savings, as the service would be paid "on-demand" for each access, instead of having to buy a complete software license.

Such "on-demand" applications are already deemed useful in the context of web services, Grid and P2P computing, which shows by the large number of academics and industry researching in these fields and the available funding programs.

In addition, the scenario assumes that services and clients will conduct business transactions, which are characterized by the exchange of payment in return for accessing the service. Most of today's Grid or P2P applications provide services for free for the participants.

However, in some Grid applications today, the service providers already receive a measurable benefit from the application, as in the distributed computation of biochemical data for a pharmaceutical company [3]. That utility gain should be divided fairly between the participants, e.g. by paying compensation for the efforts. With this fundamental motivation in the background, we expect the situation to change towards a market for services and the use of (micro-) payments.

In such a web services market, a central question is who controls the matching between service clients and service providers, as this position holds the key to manage how much utility gain every participant will receive. In general, that kind of application layer network control is conducted either centralized (by some dedicated, objective and trustworthy coordinator) or decentralized (by letting the participants negotiate directly with respective clients and suppliers).

The goal of the CATNET project is to evaluate the Catallaxy paradigm for decentralized operation of application layer networks in comparison to a Baseline centralized system.

The simulator implements two main control mechanisms for network coordination: a *Baseline* control mechanism and a *Catallactic* control mechanism. The Baseline control mechanism computes the resource allocation decision in a centralized service/resource provider. The Catallactic control mechanism has the characteristic that its resource allocation decisions are carried out by self-interested agents with only local information about the environment. Each agent has a resource discovery facility and a negotiation strategy module.

The following class types are defined in both control mechanisms:

- Client: a computer program on a certain host, which needs access to a web service to fulfill its design objectives. The Client (C) tries to access that "service" at an arbitrary location within the computer network, uses it for a defined time period, and then continues with its own program sequence. Client programs run on a connected network "resource".
- Service: an instantiation of a general application function, embodied in a computer program.
- Service Copy: one instance of the "service". The Service Copy (SC) is hosted on a "resource" computer, which provides both storage space and bandwidth for the access of the service.
- Resource: a host computer, which provides a limited number of storage space and access bandwidth for service transmission. Resources (R) are connected to each other via dedicated network connections.
- Network Connections: These connections are intended to be of equal length and thus of equal transmission time and costs.
- Node: every network element able to relay packets in the network. Every Resource is a node, but not all nodes have to be resources.

Real world applications like multimedia content distribution networks (e.g. Morpheus or AKAMAI), Grid implementations, and Peer-to-Peer systems (for instance Gnutella) are distributed applications and can be characterized in a simplified form by a number of a few common features, which inspired the design of the application layer network implemented in the simulator. Though different in many particular mechanisms, these real world applications are mapped to the following two-dimensional design space of 1) the node dynamics; and 2) the node density.

- *Node dynamics* measures the continuous availability of service-providing (resource) nodes in the network. Low dynamics mean an unchanging and constant availability; high dynamics are attributed to a network where resource nodes start up and shut down with high frequency.
- *Node density* measures the relation of service-providing (resource) nodes to the total number of network nodes. The highest density has every network node providing the described service to others; The lowest density is reached if only one resource node in the whole network exists.

In Figure 1 we illustrate the approach on how we map real world application networks in a two-dimensional space. This classification allows, mainly by means of different setup parameters of the simulation, to simulate different application layer network scenarios. The technical process of conducting the simulations is described in deliverable D2 of this project.

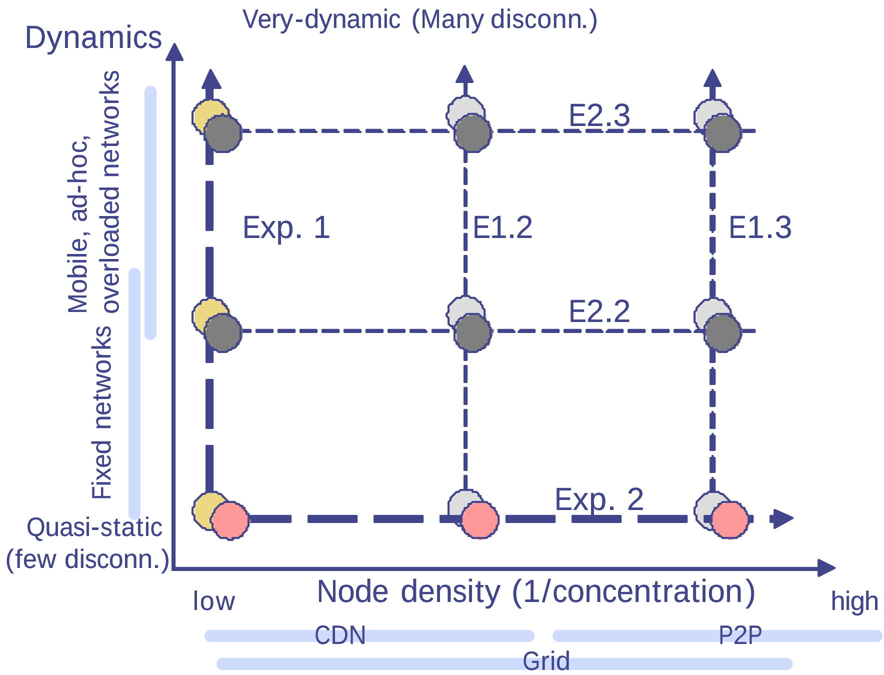

Figure 1: Mapping of real ALN to a two-dimensional design space

1.2 Scenarios

This chapter describes the different scenarios investigated by CATNET and mapped in Figure 1.

1.2.1 Quasi-static Scenario

A quasi-static application layer network can be described as having low node dynamics (quasi-static) and low node density. An application example for this are content distribution networks (CDN) such as AKAMAI. These have low node dynamics, because resource nodes are highly available and permanently connected to the fixed network. They have low node density due to the fact that resource nodes are few in terms of the total nodes which form the communications network.

1.2.2 Highly-Dynamic Scenario

P2P networks are a good example of the other extreme, highly dynamic networks. They have high node dynamics and high node density. In this case, high node dynamics is represented by the high frequency of nodes connecting and disconnecting from the P2P network. High node density is due to the fact that each peer carries out the function of a client, service provider and resource.

1.2.3 Low Node Density (and High Dynamics)

Between the extremes of CDN and P2P scenarios, one can find numerous examples of other realistic networks. The low node density scenario may correspond to web service scenarios with one or a few dedicated web servers, like in the case of Internet radio. Its high dynamics are a result of changing and unreliable connections between the network nodes.

1.2.4 High Node Density (and Low Dynamics)

This scenario combines high density (which also means spread resources) and an unchanging network. It may correspond to a very stable P2P network, an extreme case of a Grid network, or even a parallel processor computing array. In these scenarios, connections are stable and unchanging and the distances between the nodes are comparably short.

1.3 Measurement criteria

The purpose of ALN lies in the timely provision of services to clients. In doing so, ALN (Grids, P2P Networks) compete with current widespread Internet-based client/server systems. The main goal of introducing ALN has been to raise productivity, such that the same output (profit, utility) can be achieved by using less input (costs). As an example, most Grid projects aim at replacing supercomputers, by providing the same quality of output (e.g. computing the results of particle accelerator experiments) with much cheaper equipment of personal computer processors. Translated to P2P, the main application replaces a centralized database by provisioning large data files in a redundant, networked file space of cheap personal computer hard disks.

In an economic interpretation, one of the main cost drivers of current supercomputers or centralized file spaces is that the investments in hardware are geared towards peak usage, so that for the most part of its lifecycle, the hardware sits idle (processors) or empty (storage). To increase efficiency leads to higher revenues per unit. By using the same investments, the output will be increased. On the other hand, the same output can be achieved using lower investments and less hardware, e.g. abundant personal computers already connected to the Internet. The utility of ALN can thus be measured in economical as well as technical terms.

1.3.1 Social Welfare Utility (SWF)

Social welfare utility (SWF) measures how well all participants of an economic system can maximize their individual utility. Basically, utility measures the fulfilment of self-interest of the participants. The concrete equation in its easiest form sums up the individual utility profits:

$$U_{SWF} = \sum u_i \tag{1}$$

Individuals are everyone who participates in the economic environment. For an application-layer network, the individual players are the clients (who demand and pay for service access), the service provider instances –SCs– (who offer access and receive payments) and finally the network resources, who offer bandwidth and storage to the service instances. The equation thus can be written as

$$U_{SWF} = \sum u_i^C + \sum u_j^{SC} + \sum u_k^R \tag{2}$$

The individual utility of these participants computes as follows.

- The self-interest of the clients is to access a PDF conversion service at the lowest cost and/or the fastest time. Utility can thus be measured either using costs or time. In an environment where services have to be paid for access, the utility gain of clients is the difference between their *private value* (of what the access is worth) and the actually paid *transaction price*:

$$u_i^C = v_{i,p} - p \tag{3}$$

- The self-interest of the service provider is to always provide access to some Acrobat service instance, such that a minimum number of service demands have to be rejected. A valid assumption, without any application relevance of Grid would be negated, is that services can charge for their access. In this case, the service provider receives a charge price from the client. As this is business, the service provider is interested in increasing revenue by either increasing turnover (more service accesses in the same time) or profit (more profit by each service access). Each redundant web Service Copy is thus a miniature business, like a retailer's branch store. Like the clients, the service providers also have a private value for service access (*sa*) . In addition, there is private value for buying network resource access (*ra*) from the hosting node:

$$u_j^{SC} = \left(p^{sa} - v_{p,j}^{sa}\right) + \left(v_{p,j}^{ra} - p^{ra}\right) \tag{4}$$

- Last, the network provider runs resources like bandwidth, processor time or storage, which form the basis for the application. These resources incur costs, and the network provider aims to fill these costs and to make profits by increasing the usage of the resources. It is therefore in his interest that service provider and clients can easily find, match and conduct transactions. His profits can be paid either by the client (as a separate part of the service access cost) or the service provider (as a separate part of the revenue). Basically, both those variants only differ in whether the payment is transparent for the client or not. Another way to decrease costs is to optimize network parameters such as bandwidth throughput, latency and network communication flows.

$$u_k^R = p - v_{p,k} \tag{5}$$

The “maximum social welfare-criterion” (SWF) balances both costs and revenue incurred by the nodes and allows comparing different variants of the Catallaxy and Baseline implementations. It has to be noted that SWF solutions are a subset of “Pareto-efficient” ones; once the sum of the payoffs is maximized, an agent's payoff can increase only if another agent's payoff decreases.

During the experiments the social welfare is measured by a gradual procedure: All agents (Service Copies, clients and resources) keep a *private value* for the goods they buy and sell, which is updated after every self signed contract, i.e. the price of a recently closed contract directly changes the *private value* in a certain direction.

When concluding a new contract the actual price will be compared to the *private value* and the difference directly contributes to the individual social welfare. Clients – who only obtain services – can add to the SWF if the agreed price is lower than the market price and vice versa. Resources and Service Copies – which sell services – can add to the utility if the negotiated amount is higher than the individual felt market price. So, the contribution to the SWF could be either negative or positive. The amount of all these SWF-movements will be summed up and allows comparison of experiments.

SWF computes a two-time aggregation, first it aggregates within one agent, and second, it aggregates over all agents.

Regarding the aggregation within one agent, several cost types and (in the case of SCs) revenue types compute into the final individual utility gain. An increase or decrease in one cost type may be compensated in revenue or even in a counter-movement of another cost type, while the result of individual utility would stay the same. The effect of an isolated change of some experimental variable has shown not to be explanative enough.

Regarding the aggregation over the whole agent population, the validity of the SWF outcome is clearly dependant on whether the setup constitutes a zero-sum game or not. In a zero-sum game, the losses of one side are compensated by the wins of the other side.

As an example, with a private value of 8.0 money units (MU) for the buyer and 5.0 MU for the seller for a particular good, a transaction price of 6.0 MU yields a utility gain of 2.0 MU for the buyer and 1.0 MU for the seller, in total 3.0 MU for the SWF. If the compromise price is 7.0 MU, the buyer yields 1.0 and the seller 2.0, but again the SWF is 3.0 MU.

If the experiment would constitute a stringent zero-sum game, the SWF should not change at all when variables are changed, but remain constant over the population. The losses of the Clients e.g. would be fully compensated by the gains of the Service Copies.

Other measurement criteria relate to the technical performance of the network, but can also explained using economic terms. It should be noted that other similar projects, geared at analysing e.g. web proxy-caching using high-level metrics [2], have shown that the choice of criteria is somewhat arbitrary, and the results are highly dependent on the particular implementation of the experimental software. We expect this to be true also for CATNET, and thus the absolute values of these criteria should be handled with care.

1.3.2 Resource Allocation Efficiency (RAE)

Resource allocation efficiency (RAE) indicates the ratio of service demands, for which the network provides a service, to all sent service demands. In other words, it measures how many requests a client has to send until a service accepts his demand and grants access.

As the request messages (and other control messages) eat up bandwidth, a higher RAE is deemed to be better both for the individual agent and for the network as a whole.

1.3.3 Response Time (REST)

Although the access time for a service can not directly compute to a cost/revenue-based utility function, every Internet user knows from personal experience that it is an important factor.

Response Time (REST) measures the time observed by the client to access the requested service. It is influenced by the diameter of the network, the available connections and bandwidth, and especially by the necessary mechanisms to establish a working link between client and service. Effectively, access time defines a second dimension in addition to SWF, as it can not be converted directly into a cost/revenue-based economic utility function.

For comparing different networks, a lower average REST is considered to be better.

1.3.4 Bandwidth utilization/Communication Cost (CC)

This parameter measures the cost of the control messages sent in the application layer network to provide a service. In our simulator, we use different message types to fulfil the negotiation protocol. To assess the bandwidth utilization, we treat the message size as a constant which is the same for all control message types, and use a global hop counter, which computes the total number of hops needed for control messages to complete the service provision for a certain demand trace.

This means that the systemwide communication cost is dependent from the number of concurrent negotiations, the number of messages in a negotiation thread (demand trace), and the distance between the communication partners.

In the case of more concurrent negotiations, the communication cost and bandwidth utilization will naturally increase. The absolute sum of communication cost should thus be evaluated with caution, as it is only comparable if the same number of nodes are present.

In the case of a monotonic concession protocol, which is characterized by using several rounds of "propose" and "counter-propose" messages, the number of messages in a particular negotiation thread is naturally higher than in a simple "take or leave it"-protocol. A single "take or leave it" like in the Baseline approach protocol implements just two messages, one "request service" and either a "accept" or "reject" as a response.

To reach a compromise in a monotonic concession protocol, both parties have to subsequently make a concession by either lowering (seller) or raising (buyer) their offer price, until the compromise level is reached. The probability to make a concession and the height of the single concession thus influence how many steps are necessary, before any two offers cross. The protocol thus implements at least two messages, an initial "propose" and either a "accept" or "reject" as a response. However, each concession-making adds a "propose" and a counter-"propose" to this list. The number of messages exchanged in the Catallactic model is thus either equal to Baseline or 2x, 3x

or even 4x higher (depending on the amount of self-interest and greediness exhibited by the single participants).
It has to be denoted, that negotiations in Catallactic model are successive and not parallel. Thus it is possible that lower communication costs occur when the close SCs naturally answer the requests earlier and therefore get a higher chance to succeed.
Lastly, the average number of hops is dependent from the density of the nodes. As the Catallactic approach tends to flood the network with demand messages of the Clients, every Client is able to reach all Service Copies in the radius of the allowed maximum hop counter. How many SCs can be reached within that radius is a function of node density. In contrast, the result of Baseline approach depends on the location of the Master Service Copy (MSC) – the central coordinator – with regard to all clients. Given an unchanging topology, a central position of the MSC and unchanging position of the clients, the average number of hops should stay nearly constant regardless of the density.

1.4 Hypotheses

The experiments are conducted under the application of a hypotheses-based framework, which lays out how the criteria are expected to measure under different scenarios. Regarding the dimensions given in Figure 1, the following “scenarios” describe the corners of the problem space.

1.4.1 Quasi-static scenario

(H1) In a quasi-static scenario using the Catallactic model

- (H1.1) the SWF is nearly equal to the results of the Baseline model,
- (H1.2) the RAE is slightly less than in the Baseline system,
- (H1.3) the REST is slightly longer than in the Baseline system.
- (H1.4) the bandwidth utilization is slightly more than in the Baseline model,

A quasi-static application layer network can be described as having low dynamics, because resource nodes are highly available and permanently connected to the fixed network; and low node density, as those resource nodes are few in terms of the total nodes which form the communications network.

If one optimal allocation exists in the system, this scenario allows its computation. The comparatively low number of nodes prevents scalability problems, while the low dynamics allows a relatively long computation time where the result still fits the problem.

In particular, the Baseline model should be able to produce near-optimal utility maximizing results for the participant’s right from the beginning of the experiment. The adaptation mechanism of the Catallactic mechanisms increases the SWF result over time so that it approaches the same level as in the Baseline model. The gap between the result levels of both models should be determined by the running time of the experiment, as the Catallactic model starts with an inferior allocation, which increases by adaptation.

The adaptation process also leads to inferior results for RAE and REST in comparison to the Baseline model, whose results are expected not to change much once they have been computed.

The higher number of messages in a particular Catallaxy negotiation, due to the sequence of "propose" and counter-"propose" messages, leads to a slower provision time of services as well as to higher bandwidth utilization.

1.4.2 Highly dynamic scenario

(H2) In a highly dynamic scenario using the Catallactic model

- (H2.1) the SWF is higher than the results of the Baseline model,
- (H2.2) the RAE is higher than in the Baseline system,
- (H2.3) the REST is lower than in the Baseline system,
- (H2.4) the communication costs are less.

Highly dynamic networks are characterized by a high level of connection and disconnection of network links, and a high density of available resources. In such an environment, the emphasis of the resource allocation lies onto real-time computation of satisfying, rather than optimal, results.

The computation of an optimal allocation needs a time span, which is inversely proportional to the number of input information. There exists thus a certain threshold, above which the frequency of changes for a given number of network participants becomes so fast that an optimal allocation computation can not be terminated in-between. The high dynamics scenario models an environment where this threshold is exceeded. As an effect, at an initial time t_0 all environment information E_{t_0} is been used as input for the allocation function. The result $f\left(E_{t_0}\right)$, computed at termination time t_1, is then applied to an environment E_{t_1}. This time difference between both environment states leads to an inferior allocation effectiveness. The higher the dynamics are, the higher is the gap between the result and the environment.

For the SWF, this means that the utility gain which can be achieved by the participants of the system should be higher in an "anytime" algorithm like the Catallactic model than the "time sliced" Baseline model. The difference should be correlated with the environment change – the higher the change frequency, the more distinctive should the difference be.

The RAE should be higher (superior) in the Catallactic model for the same reason. RAE looks especially at open offer bids; in a highly dynamic Baseline model, those offer bids, existing in t_0, might already either have been matched in t_1 or the node is no longer available. In that case the Baseline MSC might pair a non-existing offer with a maybe also non-existing demand, which leads to an increased number of "reject" messages. In contrast, the Catallactic model should work faster, so that the non-availability of supply or demand has lower effects on the overall result.

The REST time is considered superior in Catallaxy to the Baseline model, as the continuous redesign of the network topology leads to a higher amount of unsuccessful allocations in the Baseline model until a successful allocation can be made.

The CC should be lower in the Catallactic model than in the Baseline model. The number of hops for a message in the Baseline model is determined by the distance of the client from the coordinator (MSC) and the applied routing. For a high dynamic networking, the routing changes often, so the hop number is not considered stable. In the Catallactic model, the hop number is smaller, as the radius of communication is more local. Outages in other parts of the network thus do not affect the communication of so many nodes, and so the communication cost should be lower.

1.4.3 Low node density scenario (with high dynamics)

(H3) In a low node density scenario using the Catallactic model

- (H3.1) the SWF is slightly lower than the results of the Baseline model,
- (H3.2) the RAE is less than in the Baseline system,
- (H3.3) the REST is longer than in the Baseline system,
- (H3.4) the CC are slightly higher.

In the low node density scenario, the ratio of Resources to all network nodes is small. This scenario may correspond to web service scenarios with one or a few dedicated web servers, like in the case of Internet radio. Communication paths are long from nearly every location in the network, and the non-availability of one Service Copy has a greater effect, as less redundant Service Copies are available. The network however is changing constantly, and communication links and routes are considered to dynamically appear and disappear at a higher rate.

For the SWF criteria, this means that also the scalability argument in favour of the Catallactic model does not hold, as relatively few computations have to be carried out. The Baseline approach is thus able to compute a (near) optimal result in a comparatively short time, which can not be approximated by the Catallactic approach.

For the same reason, the RAE is inferior than in the Baseline system. The number of propose messages is low, and can be answered fast by the Baseline MSC. Although the number of available Service Copies is small, each has a larger share of the totally available bandwidth to offer, and so more "accepts" than "rejects" will probably be submitted.

Due to the high concentration of SCs on only few resources, a remote client might need a long REST time to discover the requested service and would have to broadcast his request to a large number of resources. If the hosting resource is not close to the client, a negotiation will take a long period of time. In the Baseline model, once demand and supply have been received by the MSC, an allocation can easily be calculated. However, the hypothesis concerning REST has not been formulated in the project proposal and is mentioned here for reasons of completeness.

The communication costs will be higher in the Catallactic approach, as the average number of hops is dependent from the density of the nodes. As the Catallactic approach tends to flood the network with demand messages of the Clients, every Client is able to reach all Service Copies in the radius of the allowed maximum hop counter. In a low node density scenario, the number of reachable SCs is thus low. In contrast, the result of Baseline approach depends on the location of the MSC with regard to all clients. Given an unchanging topology, a central position of the MSC and unchanging position of the clients, the average number of hops should stay nearly constant regardless of the density.

1.4.4 High node density scenario (with low dynamics)

(H4) In a high node density scenario using the Catallactic model

- (H4.1) the SWF is equal or better than the results of the Baseline model,
- (H4.2) the RAE is slightly less than in the Baseline system,
- (H4.3) the REST is higher than in the Baseline system,
- (H4.4) the CC are lower.

In a high node density scenario, the communication paths in the Catallaxy model are shorter and there are more SCs within a Clients hop counter range. This leads to higher competition in the market, as more alternatives are available. At the same time, the network has a high reliability of providing communication links. This scenario may apply to clustered processors in a local network or parallel supercomputing.

The SWF result of the Catallactic model is considered to be superior to the Baseline model, because a greater choice in transaction partners leads to the ability for each Client to individually select one SC who maximizes the individual utility goal. The higher the dynamics are in addition, the lower the possibility that the slower Baseline approach would also come up with the same pairing. It is thus expected that the individual utility gain from each transaction is greater, even if not all requests can be filled.

This hints to the development of the RAE criteria. As the competition increases, numerous negotiations are done in parallel and even the local partner selection process of the single agents becomes slow in adapting the changes in the environment. As a result, many negotiation attempts are tried until a particular negotiation succeeds. This leads to a higher number of “rejects”, even though the result from the individual “accept” is better.

For the same reason the REST is considered to be higher in the Catallactic model. A high number of negotiations will lead to more efficient allocations but will therefore have to be paid by longer response times.

The communication costs in the Catallactic model are lower, as the messages do not have to travel far. Within each hop range diameter are sufficient competing SCs and Clients, so that satisfying results can be reached.

2. Comparison of experimental results

2.1 Experimental setup

We use the nomenclature given in Table 1 to identify the experiments. All experiment data ends on a certain experiment code, defined by the appendix *_xy*, where *x* is the level of dynamics, and *y* is the level of node density

Table 1: Nomenclature for experiment description

Experiment code	Description
_00	Low node dynamics, low node density
_10	medium node dynamics, low node density
_20	High node dynamics, low node density
_01	low node dynamics, medium node density
_11	Medium node dynamics, medium node density
_21	high node dynamics, medium node density
_02	Low node dynamics, high node density
_12	medium node dynamics, high node density
_22	High node dynamics, high node density

Graphically, these scenarios are positioned in the experiment space as shown in the following figure.

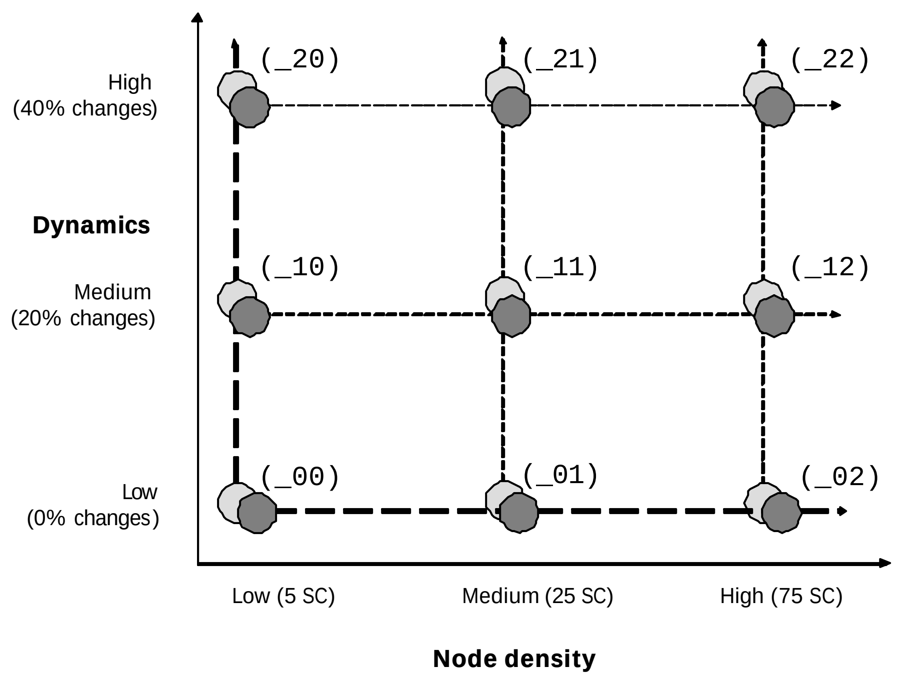

Figure 2: Experiment Space

The most common setup is shown in the following figures. It depicts how the results of the experiments can be configured, distinguished and archived in an excel-sheet. The technical description has been largely given in the Deliverable D2 of the CATNET project. The topology of the experiment is usually unchanged and looks like Figure 3.

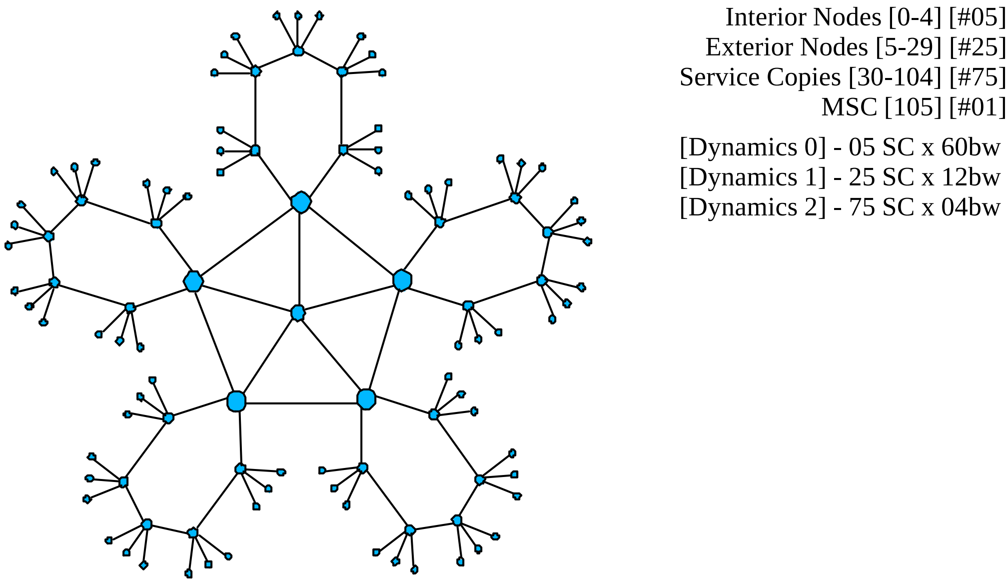

Figure 3: Network Topology for Experiments

The network consists of one inner ring and several branching rings. The number of nodes in total is kept constant at 106 nodes, on which 75 Client instances are distributed (see Figure 4). In the basic Baseline models, the MSC issues a new allocation every 50ms. In the sensitivity analysis described later, we will change this MSC Allocation time to investigate its effect on the network performance. The 75 clients issue 2000 demands for service access during a given runtime of experiments of 100 seconds. This means that every 50 milliseconds a new demand is posted in the network.

The different dynamics for the _0x, _1x and _2x type of scenarios are defined by making 30% of all Service Copies initially available. During the runtime this availability will be changed every 200ms with a probability of 0% for the static (_0x) case, 20% for the medium (_1x) and 40% for the high (_2x) dynamic case.

The different densities for the _x0, _x1 and _x2 scenarios are defined by how many Service Copies are available in the network so that the total amount of service access (measured in slotted bandwidth) for the Clients remains constant. In the low density experiments (_x0), the total number of 300 slots is distributed by 5 SCs, each providing 60 slots at once. In medium density (_x1), 25 SCs provide 12 slots each, and in the high density scenario (_x2), 75 SCs provide 4 slots each for a total of 300 slots.

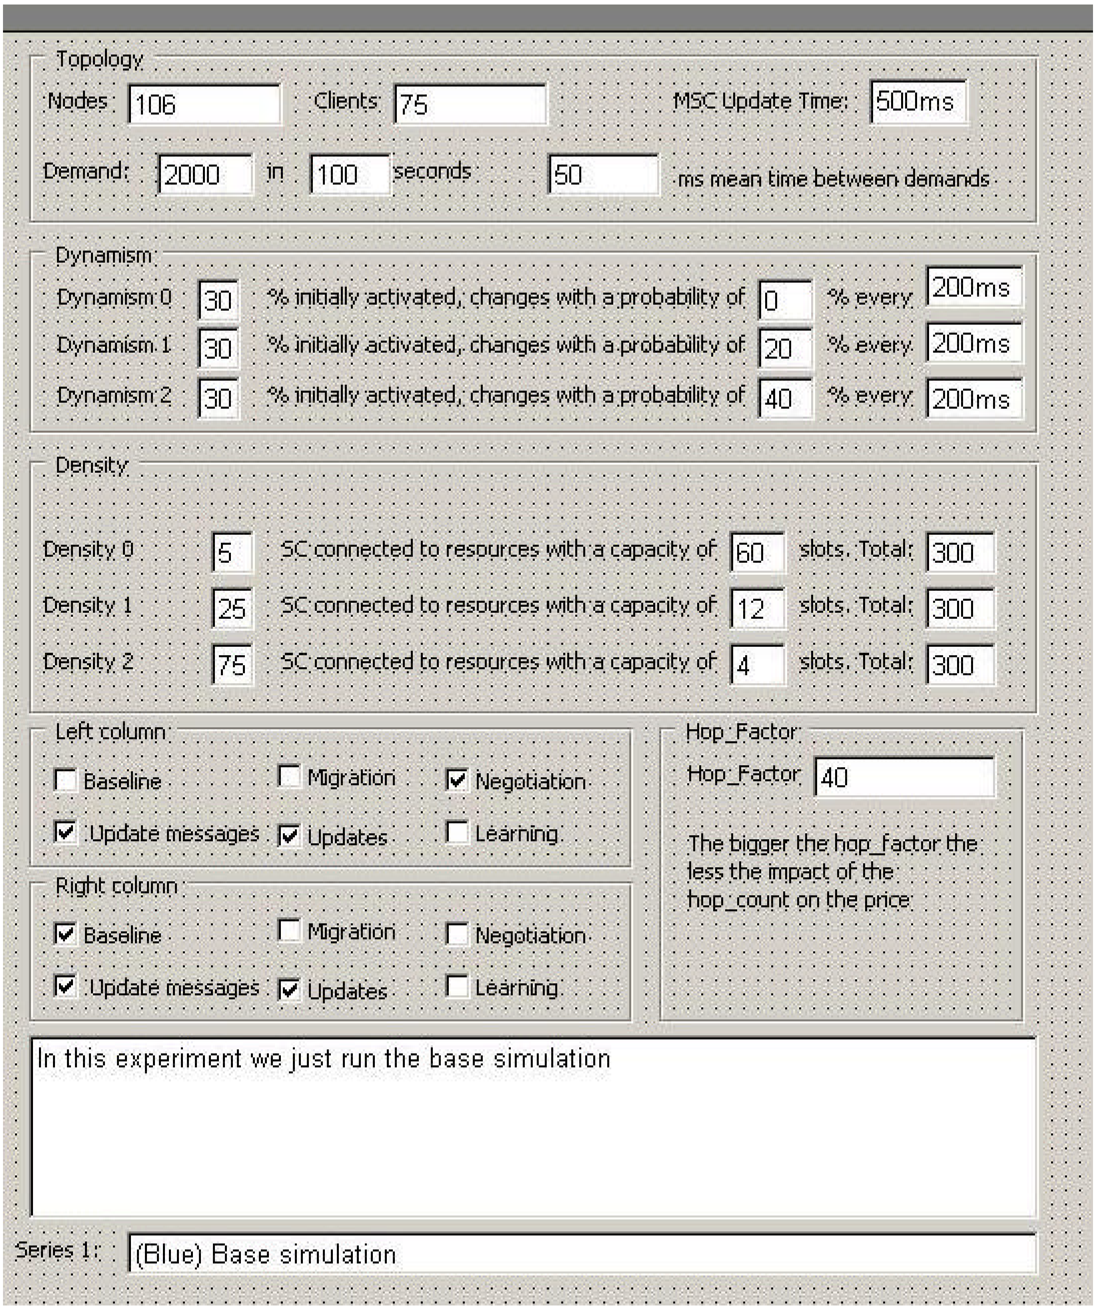

Figure 4: Variable Setup for Experiments

2.2 Quasi-static scenario (_00)

This type of scenario resembles content distribution applications such as AKAMAI which have low node dynamics, so that resource nodes are highly available and permanently connected to the network, and low node density so that resource nodes are few in terms of the total nodes which form the application layer network. The graphics displayed in this paragraph depict the difference between Catallactic and Baseline and have been normalized with respect to Baseline, e.g. for the SWF value:

$$SWF\% = \frac{(SWF_C - SWF_B)}{SWF_B} * 100 \qquad (6)$$

This gives a fast visual representation of the importance of a change with respect to Baseline. As the exact value of each parameter is not relevant (it depends on the concrete experiment we have used), the resulting information is not changed.

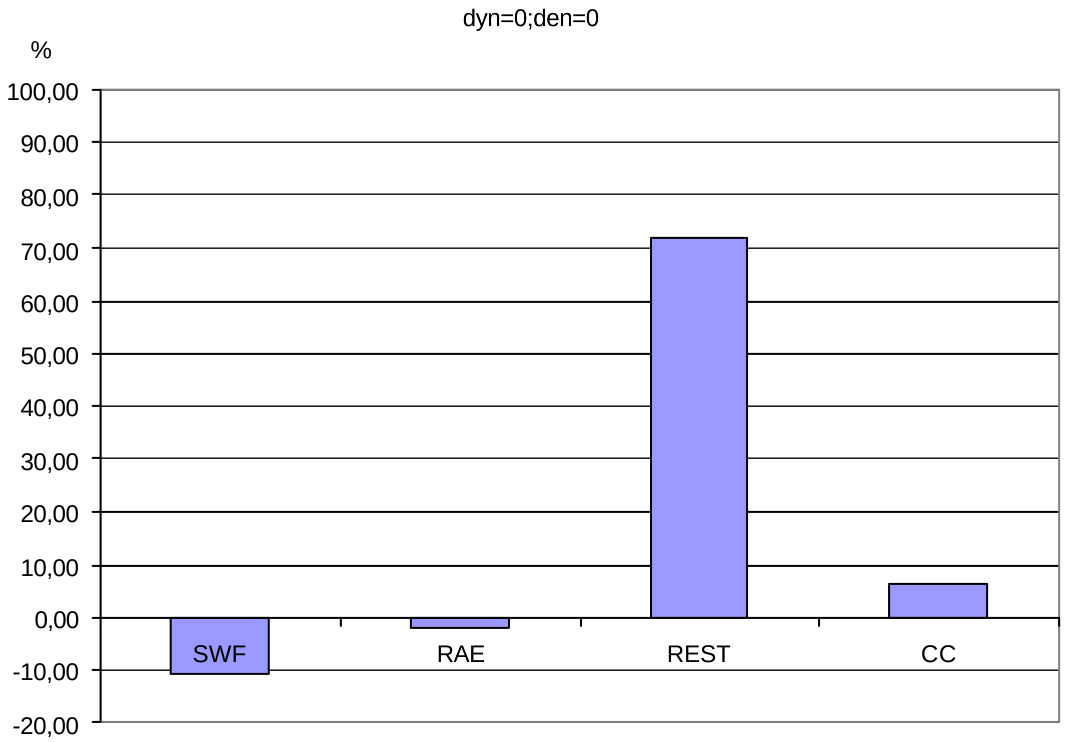

Figure 5: Results of the Quasi-static scenario

The results of the SWF criterion show that the Baseline model outperforms the Catallactic approach by more than 10%. This is near hypothesis H1.1 and shows that Baseline coordination is slightly better for this scenario with respect to SWF.

RAE is nearly equal in both scenarios. Hypothesis H1.2 expected that RAE should be a little less in the Catallactic scenario than in Baseline, this could be shown.

The response time (REST) for the Catallactic model is nearly double as high as for Baseline. This confirms the tendency of hypothesis H1.3, but falsifies the expected order of magnitude.

The communication costs CC have not been part of the original hypothesis set. However, combined with the performance of the other criteria, the development of this factor is not unexpected. The communication cost are higher in the Catallactic case, which means that more messages are exchanged, which in turn results in more bandwidth utilization. H1.4 can be confirmed.

2.3 Highly-dynamic scenario (_22)

Peer-to-Peer (P2P) networks can be described as networks with high node dynamics and high node density. In this case, high node dynamics is given by the high level of connection and disconnection found in P2P networks. High node density is illustrated by the fact that each peer carries out the function of a client, service provider and resource.

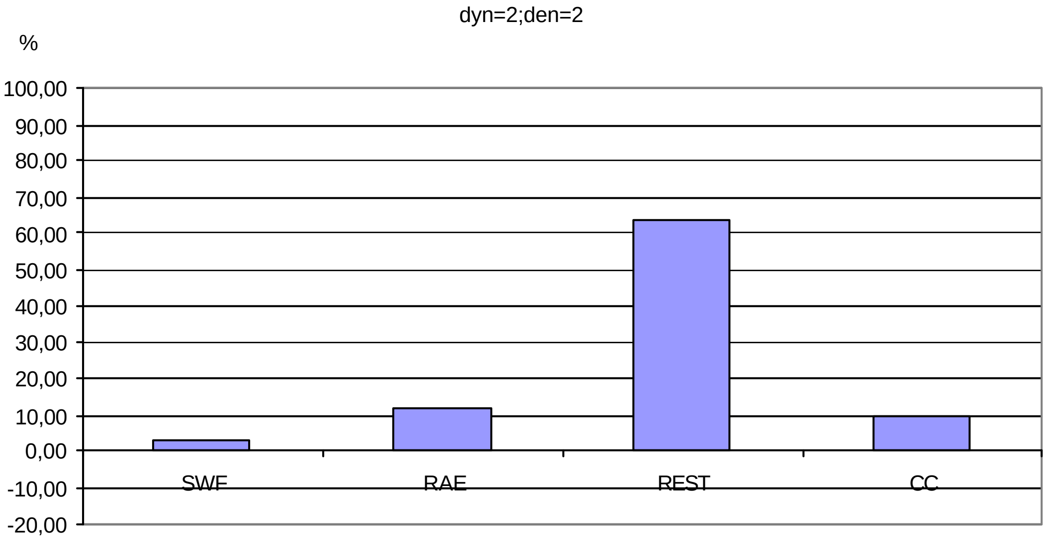

Figure 6: Results of a highly-dynamic scenario

The SWF value is nearly equal for both Catallaxy and the Baseline case, which falsifies H2.1. Our expectation that Catallaxy would perform significantly better can not be confirmed.

The resource allocation efficiency (RAE) is higher in the Catallactic model, which confirms H2.2.

The response time is significantly longer in the Catallactic model, which falsifies H2.3. This is considered to be an effect of the continuous change of topology in the network which leads to a higher number of negotiations. These negotiations take longer, the time between issuing a demand and receiving a positive allocation result will thus take longer.

The communication costs (CC) in the Catallactic model are still higher as in the Baseline model, which falsifies H2.4. An explanation may be that the number of control messages needed to achieve a service is significantly higher for Catallaxy as expected previously.

2.4 Scenario low node density (20)

Low density (which also means high resource concentration) may correspond to web service scenarios with one or a few web servers. At the same time, the dynamics of the network are quite high, which reflects to unreliable long distance connections between the nodes.

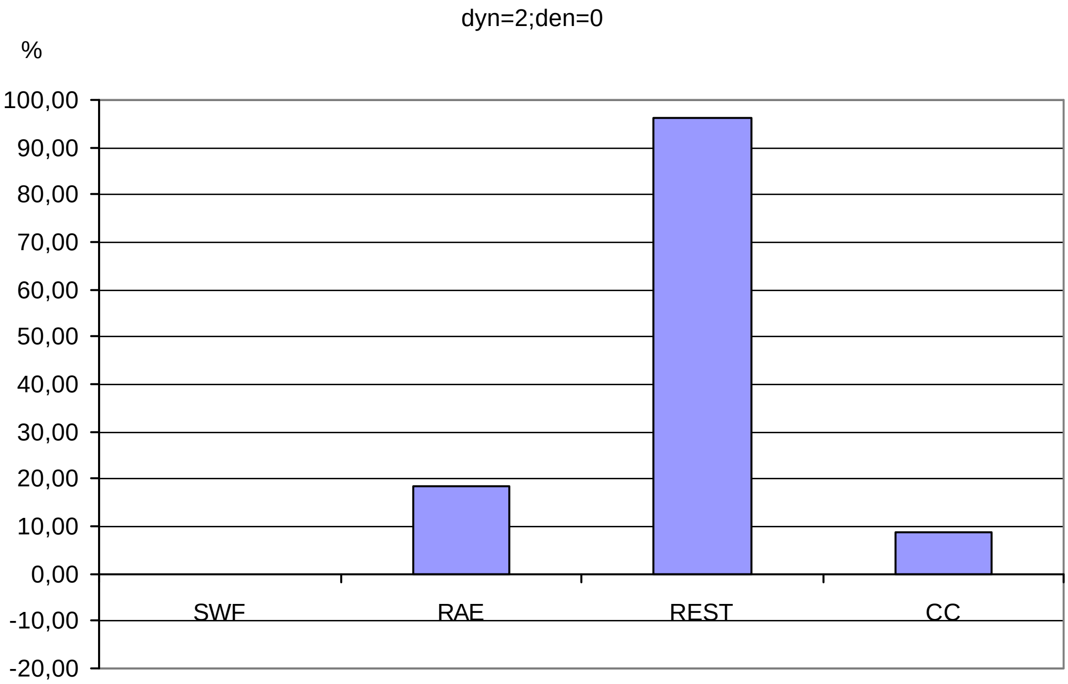

Figure 7: Results of low-node density experiment

The hypothesis H3.1 expects the social welfare utility (SWF) of the Catallactic model to be inferior to the Baseline model. However, both Catallaxy and Baseline are nearly equal. Hypothesis 3.1 is thus rejected.

The resource allocation efficiency (RAE) is surprisingly also better for the Catallactic model than for the Baseline system. H3.2 is rejected by this.

The response time (REST) is longer for the Catallactic model, confirming H3.3 by this.

The communication cost (CC) are higher for the Catallactic model, confirming H3.4.

In total, this scenario poses the most challenging task for interpretation, as nearly all hypotheses were rejected.

2.5 Scenario high node density (_02)

A high density scenario may correspond to a P2P network, an extreme case of Grid, clustered processors in a local network or parallel supercomputing. Overall, the dynamics of the scenario are low, so the links between the nodes are stable. Once paths and routes have been found, they need not change.

The Catallactic model achieves a slightly inferior SWF to the Baseline model, thus neither really confirming nor rejecting H4.1 which expected an “equal or better” mark. The resource allocation efficiency (RAE) is slightly worse in Catallaxy, which confirms H4.2 which expected the Baseline model to be superior. The response time (REST) is higher in the Catallactic model, which confirms H4.3. The communication cost are also higher, which falsifies H4.4.

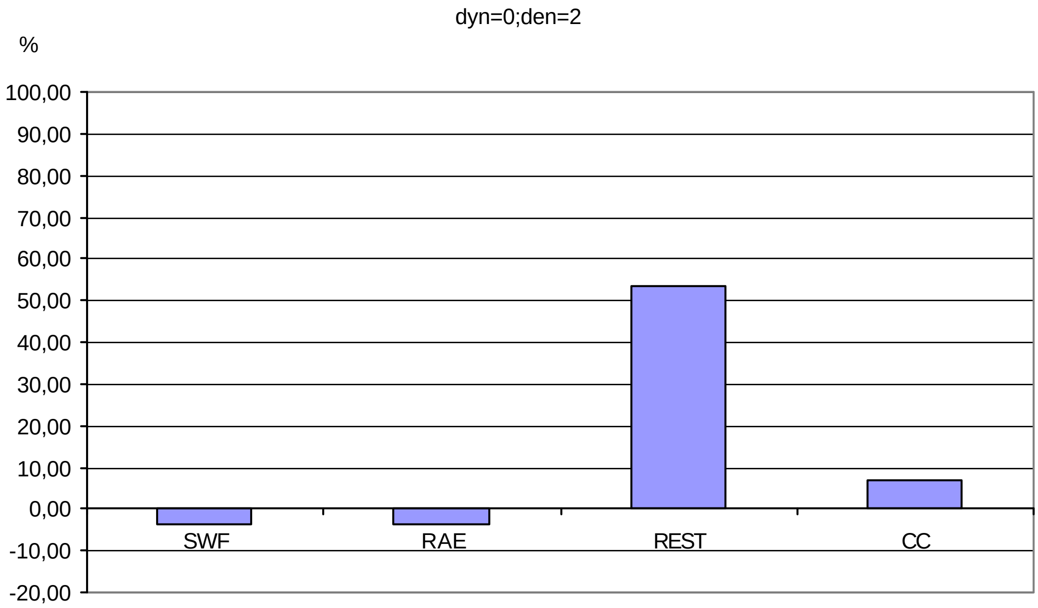

Figure 8: Results of the high density scenario

2.6 Experiment results by criterion

To compare the different experiments with each other, the following section shows the same experiments and values (improvement of Catallactic over Baseline in %) as above, but organized in a 3D figure for each criterion with results from all combinations of 3 values of dynamics and 3 values of density.

2.6.1 Social Welfare Utility (SWF)

The first factor to be looked at is social welfare utility (SWF). In the result graphics, the Baseline and Catallactic model for each experimental setup are usually displayed as pair of bars for each of the 9 scenarios, where the left bar shows the Catallaxy result and the right bar the Baseline results for whatever criterion (SWF, REST, RAE, CC) is displayed (see Figure 9). If several rows exist, a sensitivity analysis of changing variable settings can be made.

The X-axis of every figure shows the 18 scenarios which are compared to each other. The picture shows 9 pairs of bars, for each spot depicted earlier in Figure 1. The left bar of each pair is the result of the Catallactic model, the right bar is the result of the Baseline model.

The Y-axis shows the absolute measure of the current variable. It should be noted that the absolute value is dependent on several factors including the topology and setup of the experiment; in most interpretations of the experiment results, we will refer to the relative value or the relation between Catallactic and Baseline result.

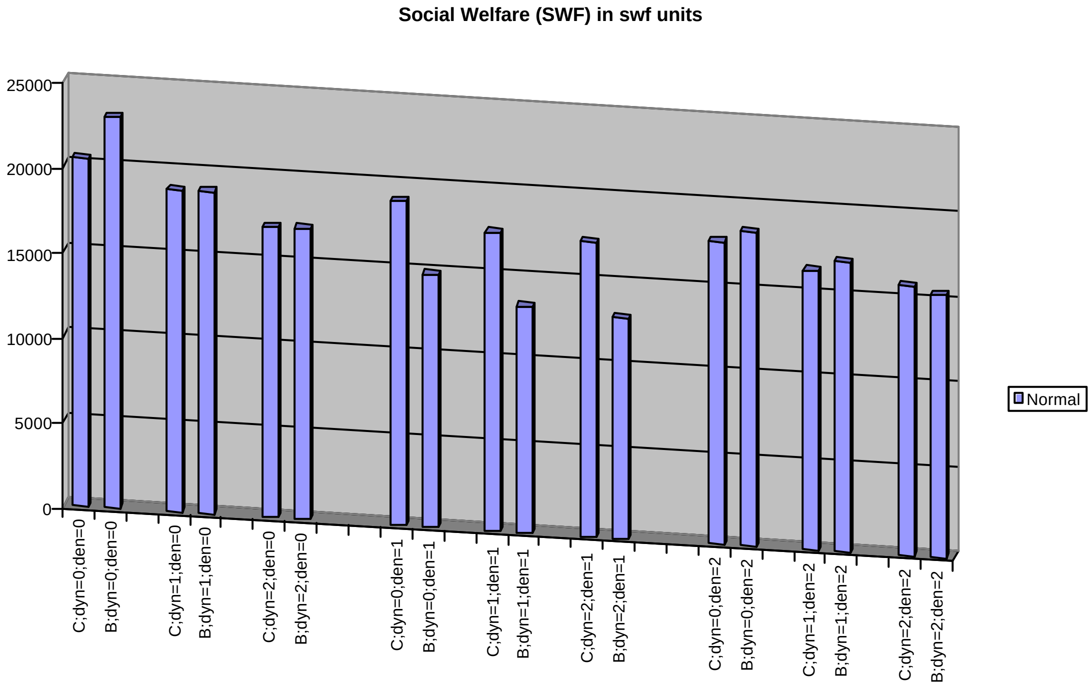

Figure 9: SWF development in different experiments

The three leftmost pairs of bars show the development of SWF from left to right under increasing dynamics, while density is kept low (scenarios _00, _10, _20). In the low dynamics/low density scenario, Baseline outperforms Catallaxy. When increasing dynamics, both Catallaxy and Baseline SWF slowly decrease with the Baseline SWF losing comparably more ground. The "break even" between both developments seems to lie near the high dynamics scenario.

This development is mirrored when changing dynamics in the medium and high density cases. Again, both Baseline and Catallaxy SWF decrease with increasing dynamics; in the high density regime, Baseline outperforms Catallaxy in the low and medium dynamics, which reverses in the high dynamics scenario.

In the middle density regime, Catallaxy SWF is always better. It also seems that the slope of change of both SWF is lower than for the other regimes.

If we look at the variation of node density under fixed dynamics, we have to compare only the leftmost pair from each triple for low dynamics with increasing density. The Catallactic model results decrease with growing density, while the Baseline SWF shows a minimum in the medium density regime.

The corresponding development for high dynamics and varying density is shown in the rightmost pairs of bars in Figure 9. Even for high dynamics, the Catallactic and Baseline show the same progress as for low dynamics. Overall, this contradicts the underlying hypotheses that Catallaxy will perform better if the environment gets more complex, and Baseline performance decreases due to growing inconsistencies and scalability problems. What holds, is that the ratio of Catallaxy vs. Baseline becomes favorable for Catallaxy with increasing dynamics.

2.6.2 Resource Allocation Efficiency (RAE)

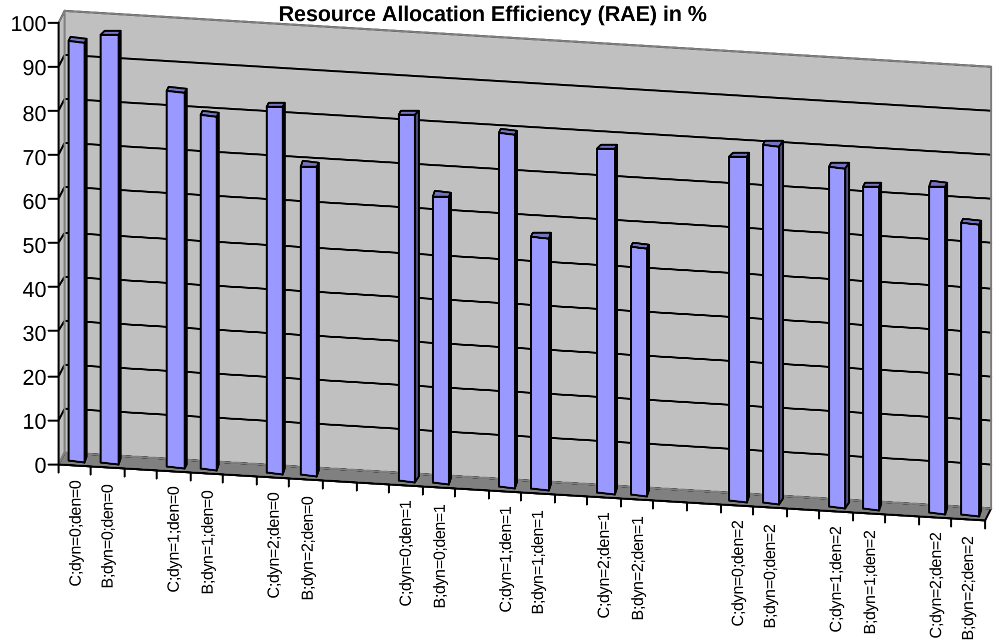

Figure 10: Comparison of RAE development in basic experiment

Regarding the resource allocation efficiency (RAE), which determines the ratio of filled demands to all sent requests of the Clients, it becomes soon clear that both models work best under low dynamics and low density.

The Baseline model achieves to match nearly 100% of all requests in the quasi-static scenario, with Catallactic model closely behind. However, as dynamics increases, the Baseline model soon loses comparably more performance than the Catallactic model.

This relative stability of results is mirrored for Catallaxy also in the medium and high density regimes. Under high density, the decrease of RAE results for Baseline is lower than in the low density regime.

Varying density while keeping dynamics steady, RAE decreases for both models similarly. For Baseline, however, the medium density regime shows a very low performance with stable results for all types of network dynamics. It is difficult to interpret this particular outcome.

Overall, the development of RAE is in the range expected by the hypotheses. This indicates, in the interpretation of the experimentators, that the hypotheses were soundly formulated, and gives a reference point for the other criteria.

2.6.3 Response Time (REST)

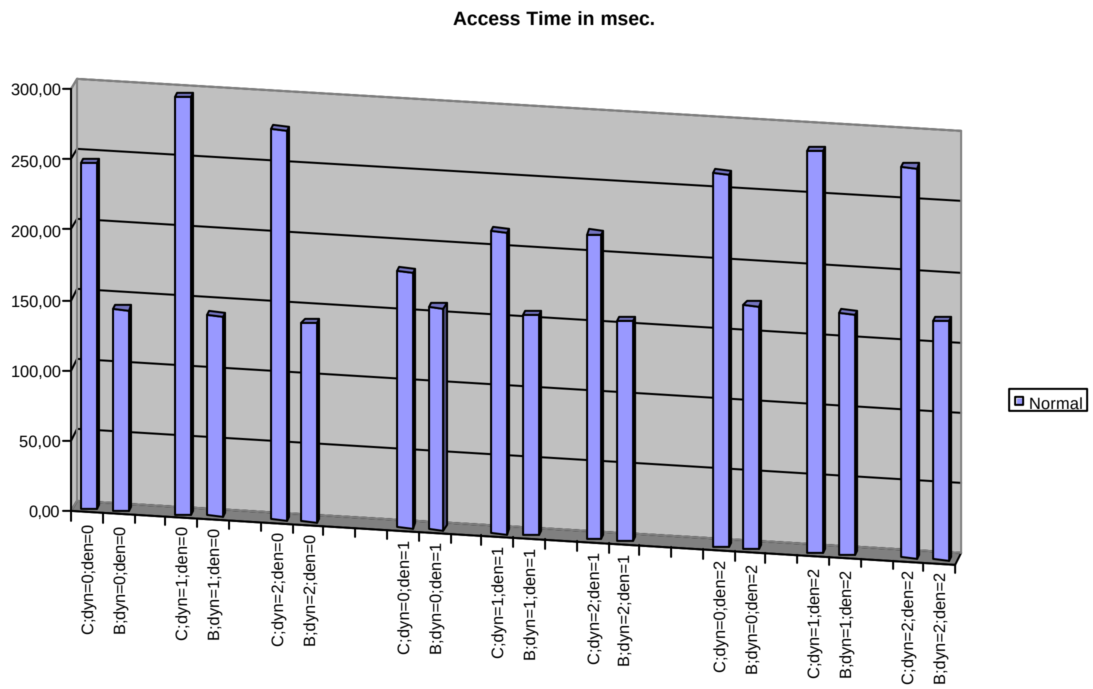

Figure 11: Response time development in basic experiments

The response time (or access time) criterion (REST) measures the time between issuing a service request and finally receiving that service. Adverse to the other criteria, a higher REST number indicates an inferior performance, as it shows a longer time between requesting and receiving a service.

Over all three density regimes, the Baseline REST increases from left to (low density) to right (high density), leading to longer access times. However, the change of the dynamics dimensions in a particular regime shows no effect. For Baseline, the REST value seems to depend on the density, but not on the dynamics.

This finding is not mirrored by the development of the Catallactic REST, which shows significant differences between and within the different density regimes. The medium density regime shows a significant faster access time than both high and low density regime, which leads to the suspicion that an optimum for REST exists (in the particular setup and topology), which depends on the density dimension only. Different from Baseline, however, a change in dynamics has a measurable effect on the outcome. In the low density scenario, the medium dynamics show the worst REST performance of all. This is mirrored in the high density scenario, where medium dynamics also leads to inferior REST time. In the medium density scenario, it seems that this artifact is not as significant.

From the results in Figure 11, it shows that changing dynamics does nearly not affect the results of the Baseline model, while the Catallaxy shows a lower performance for the medium dynamics regimes.

With regard to the hypotheses, the response time criterion was expected to be in favor of Catallaxy for a lower density and Baseline for a higher density. This statement can not be confirmed, as the REST is inferior for Catallaxy in all scenarios. In reverse, the development of the REST parameter seems to be dependent from other factors than those considered in the hypotheses. Overall, the REST development leaves a mixed bag of statements about the hypotheses.

2.6.4 Communication Cost (CC)

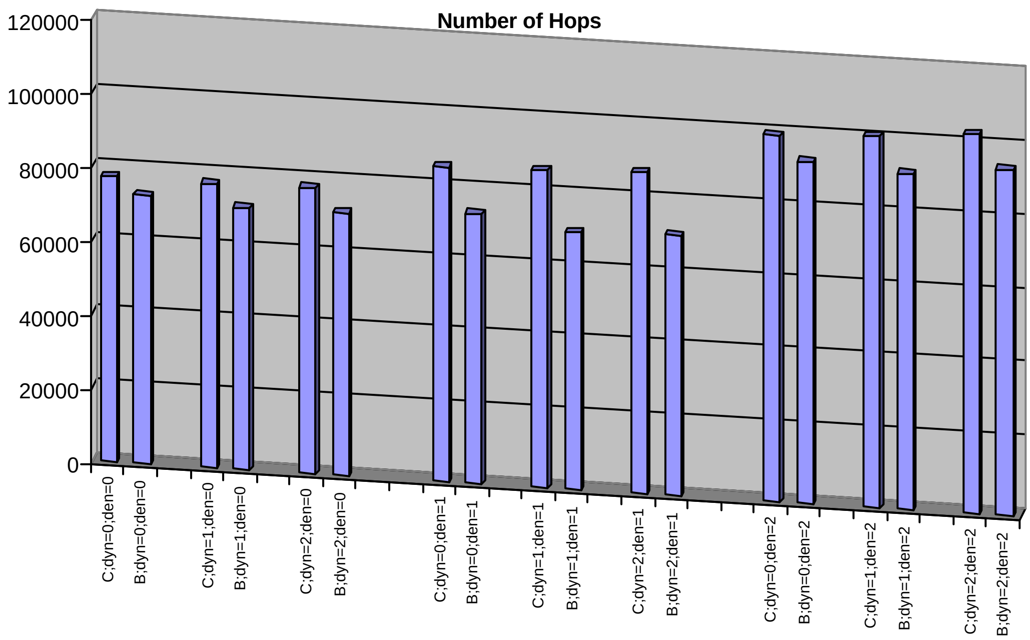

Figure 12: Development of CC in basic experiments

Like for the response time, the communication costs (CC) results are inferior if the values are higher. Figure 12 thus shows an overall picture of defeat for the Catallactic model.

As expected by the experimental setup, the communication costs raise with the availability of a growing number of nodes. It comes thus not as a surprise that growing density leads to increasing costs for both models.

The variation of dynamics, however, does not affect the results much, which remain steady if density is constant.

Overall, the CC performance matches the expectations and hypotheses. It should be noted, that the medium density experiments, which showed some odd numbers for the other criteria, seem here to be perfectly in line.

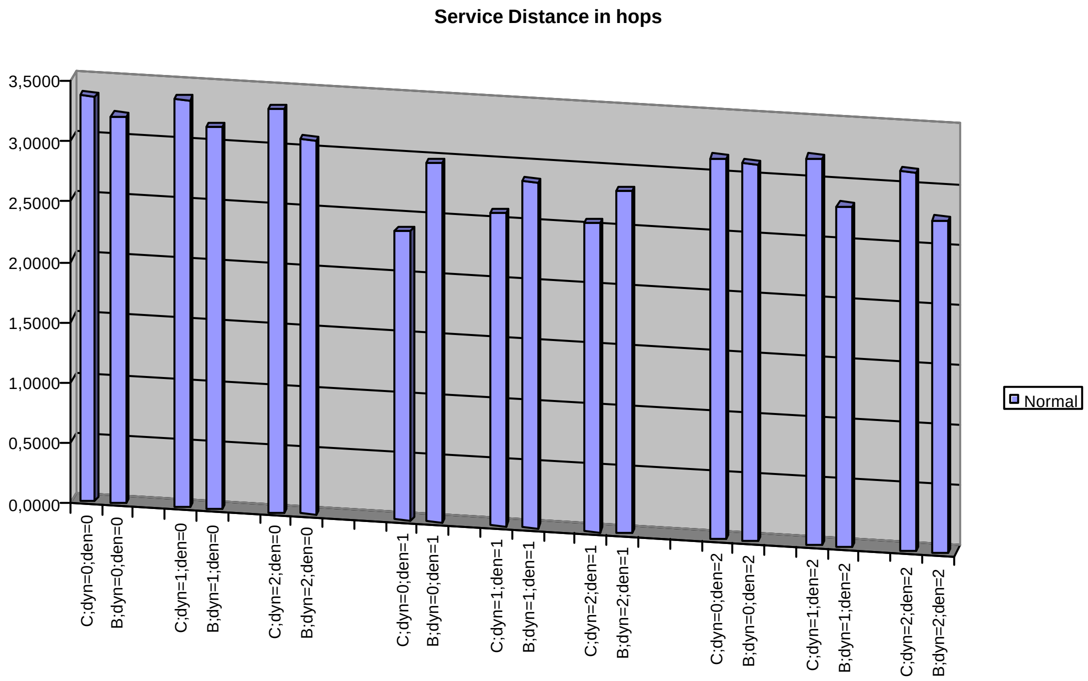

Figure 13: Development of average service distance

As a further result, which is related to the communications cost, we have measured the average distance between matched Clients and Service Copies, in network hops.

Overall, the figure shows that the matching conducted by the Baseline model does not take the distance between Clients and Service Copies into account. For all experiments, the Catallactic model achieves a shorter distance between access demand and access provider.

In the low density case, this distance is naturally at a maximum, as the Service Copies are more spread in the network. With growing density, the resources are closer together and it is possible to find a nearer provider for the same local demand. This statistical phenomenon is shown quite clearly if one compares the Baseline results for low and high density. Without a change in the MSC matching strategy, the average service distance is lower for the high density regime.

As the Catallactic model routes its demands locally through the network, the Clients will very often begin negotiating with the nearest Service Copy and possibly reach an agreement. From this behavior follows a tendency to have shorter distances between Clients and Service Copies.

As with some of the other criteria, the medium density experiments show a particular result in that the average distance is at a minimum for both models. This may be caused by a favorable relation of the number of Clients to the number of Service Copies, exploited by the particular mechanisms in the experimental setup.

Overall, the fundamental assumptions of the hypotheses hold and lead to an expected development of the average service distance. However, the particular differences between the varying dynamics and the varying density illustrate the necessity to further investigate into the relations between the different criteria and the experimental setup. This can be done by doing some sensitivity experiments, as shown in the next section.

3. Sensitivity experiments

The basic experiments described in the previous chapter have been conducted to show the fundamental differences between the Catallactic and the Baseline models. The hypotheses described in the beginning of this report could not be thoroughly confirmed or rejected, as interdependencies between criteria and experimental setup variables (which have not been accounted for) modify the result and make interpretation difficult.

In this section, we conduct some sensitivity experiments in order to change experimental setup variables, like the MSC Update time, the demand frequency, or whether Service Copies can relocate (migrate) from one resource to another. Figure 4 on page 17 has shown which setup variables are available and might be changed.

3.1 Effects on changing demand frequency

The demand frequency describes the time span between two request messages launched by each client. This continuous message sending can be considered as a steady demand pulse, whose frequency influences the measured parameters by increasing the charge of the system. The following graphics show the behaviour of the simulator when changing the demand intensity from 25ms over 50 ms and 100ms to 200ms.

The Z-axis shows in different rows the results of a change in a particular variable, here "demand frequency". Depending on the actual experiment, several rows are possible. Here again, the absolute value may be of lower significance than the relation to other rows of the same scenario or the respective Catallactic or Baseline counterpart. Sometimes, an experiment could not be played fully through due to technical reasons. In that case, the respective bars for that particular variable setting are not shown.

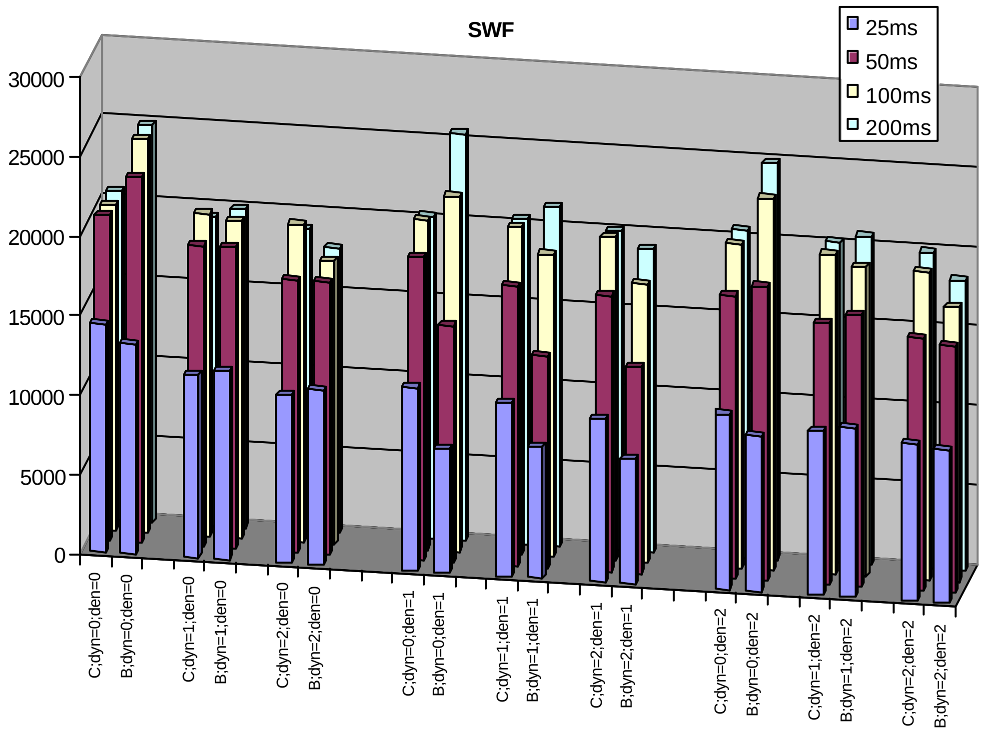

Figure 14: Effects of demand frequency variation on SWF

A higher demand frequency leads to a higher system load and to congestions, and thus to less closing of contracts. SWF increases with a slower demand rate; in all columns displayed in Figure 14, the SWF value is highest when the demand frequency is at its lowest, 200ms. This effect can be well observed in low density networks, where congestions occur because of limited network resources. A wider distribution/higher density leads to a relaxation and better SWF results which compensate the decrease of SWF. Increasing dynamics goes along with decreasing SWF, regardless of whether it is Baseline or Catallactic model.

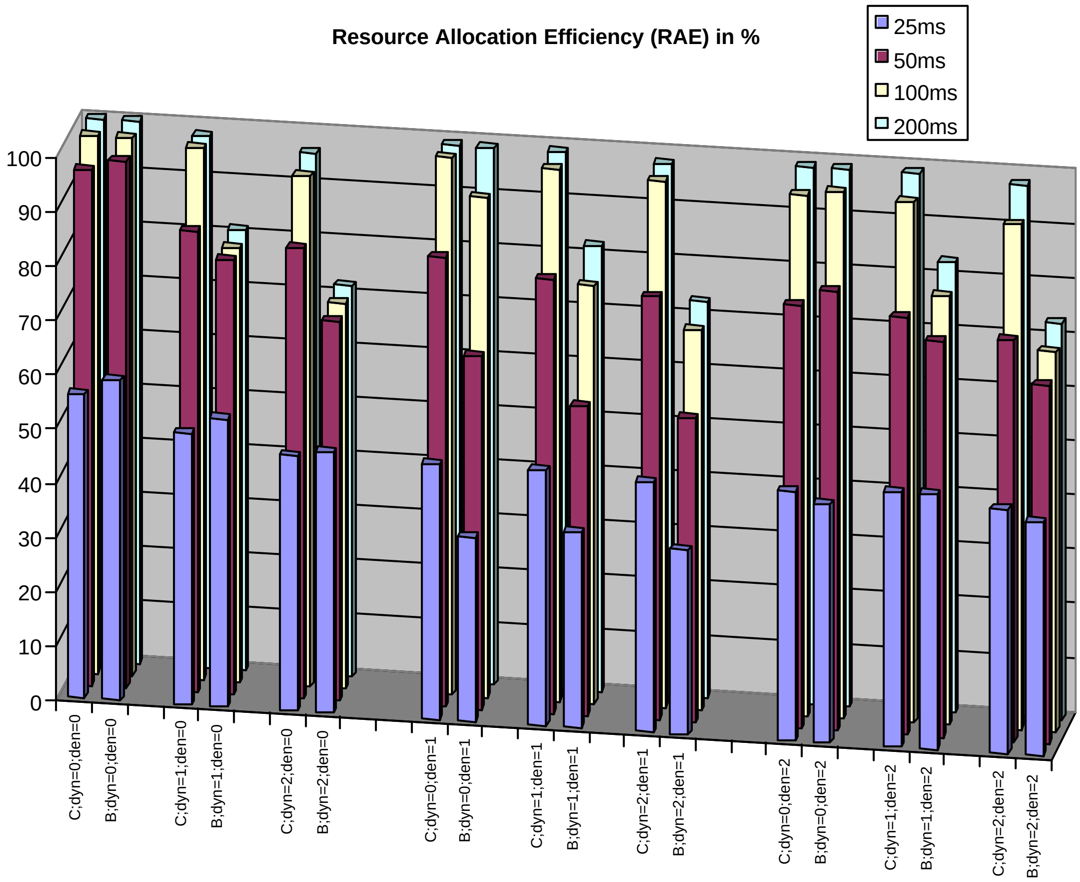

Figure 15: Effects of demand frequency variation on RAE

Considering the RAE, the results are mostly the same than in the SWF argumentation above. 100% RAE can be achieved for the Catallactic model with a 200ms pulse even in high density. Complete saturation of demand in Baseline can even be attained in a low dynamics/high density scenario. Increasing the demand frequency implies a high charge to the system and in case of 25ms pulses reduces the Resource Allocation Efficiency to nearly 50% in all scenarios.

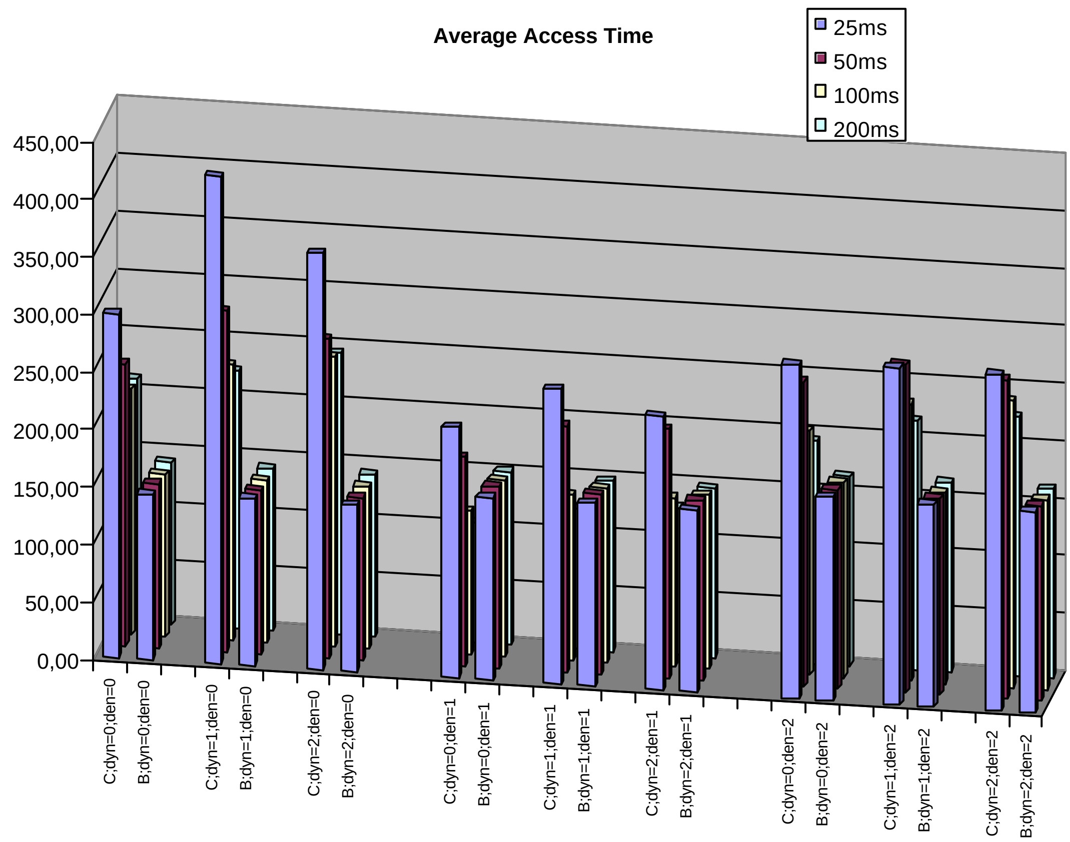

Figure 16: Effects of demand frequency on access time

For the REST criteria in general, a lower demand frequency leads to faster access times. This effect could be caused by an increasing synchronization of the frequency demand generation with the negotiations which lead to successful service provisioning. In other words, the large number of open demands leads to longer queues for processing them in the Service Copies; as the queues get shorter with decreasing demand frequency, the overall processing gets faster. The REST parameter for Baseline keeps always the same. This is a consequence of the steady MSC Allocation Time (50ms) which starts after receiving the first request for a certain item by the MSC and thus is more relevant to the REST.

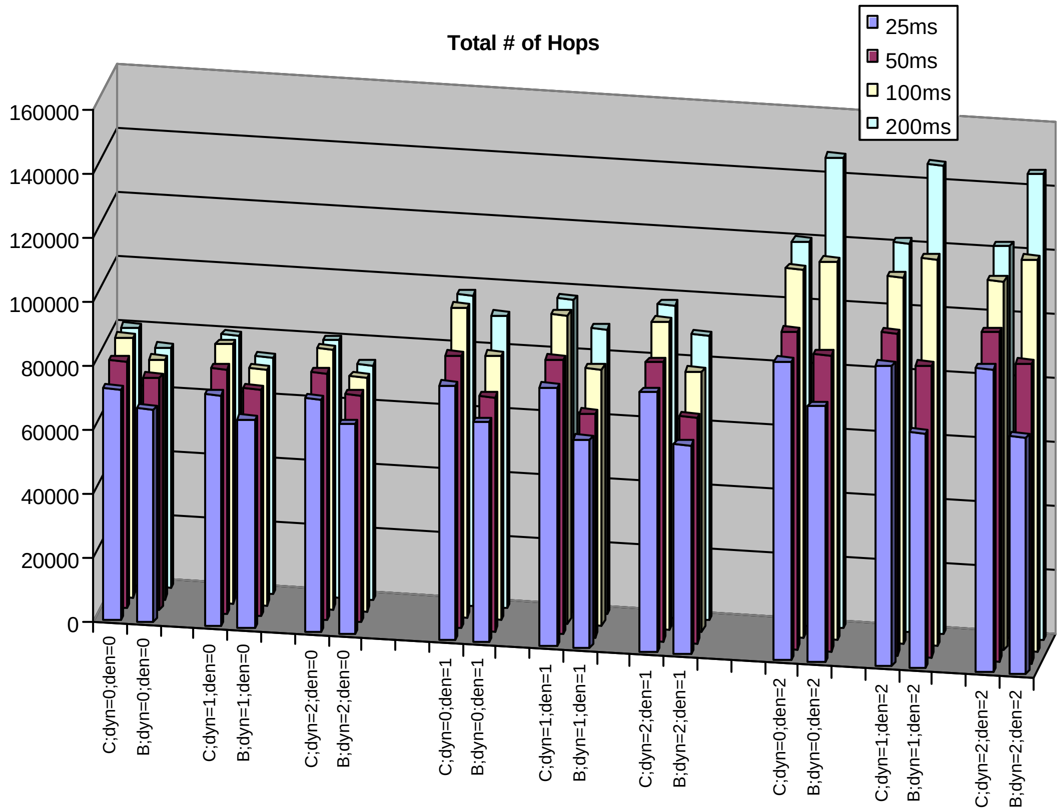

Figure 17: Effects of demand frequency on the hop count

A particular development can be shown when looking at the message hop count. Here, the number of hops increases in correlation with the increase in density. Changing dynamics, however almost does not change the amount of hops. This result correlates with that of the basic experiments as shown in Figure 12.

Regarding the sensitivity analysis, the increasing demand frequency has an adverse effect on the hop count. Both for Baseline and Catallactic models, the sum of the message distance decreases when the frequency gets higher. This could be explained by the low quantity of negotiations initiated because of rising queue lengths at resources and Service Copies due to the fact that negotiations take place sequential. This approach is sustained by the immense decrease of RAE at higher demand pulses.

Summing up, the demand frequency parameter has some great influences on the experimentation results and has to be selected with strong regard.

3.2 Effects of changed MSC Update time

To keep the MSC informed about the actual state of the Service Copies, an update message system was introduced for investigational reasons. The Service Copies will have to send an "I am alive" message to the MSC after a given period of time. This time period can be analyzed in relation to the experimental results. The following figures will display the behavior of SWF, RAE, REST and CC. As update time exclusively affects Baseline model, only the right bars are considered for the analysis.

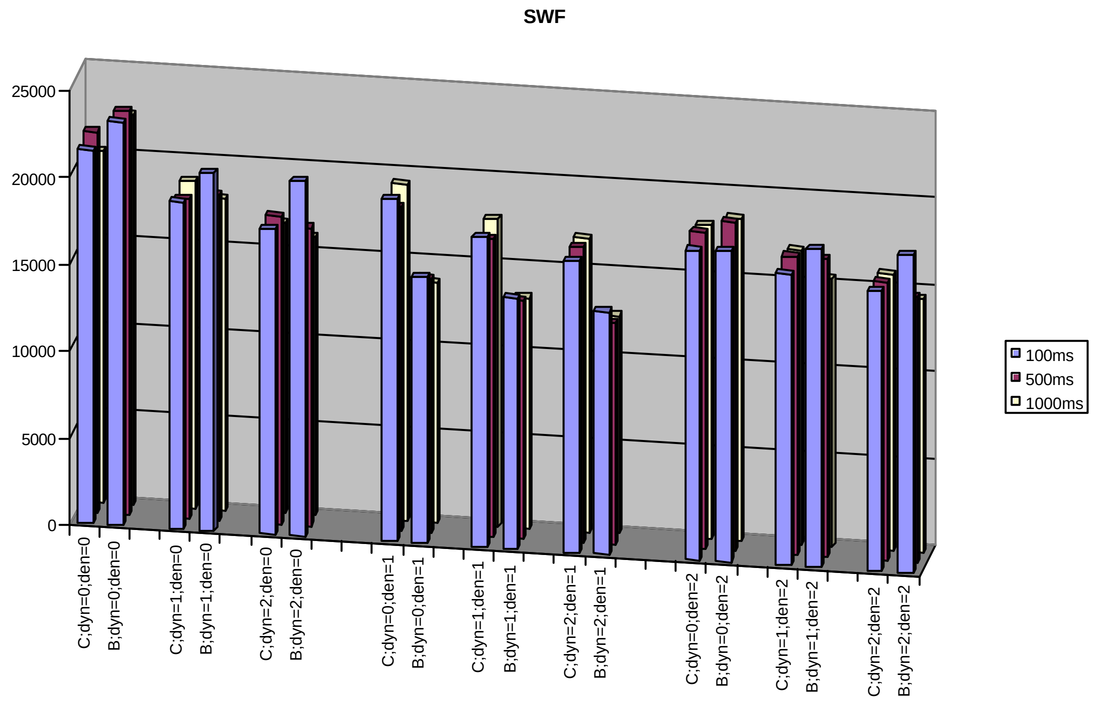

Figure 18: Effects of MSC Update Time on the SWF

In conjunction with higher update message frequencies, the MSC keeps informed up-to-date and always has a close impression of the "reality". This leads to better results and less allocation failures.

It can be observed that the results show an impact on the SWF. As expected, in quasi-static scenario, the change of MSC Update time has nearly no effect on SWF. Amazingly SWF increases slightly when lowering the frequency. This cannot be explained and is considered to be in the range of usual variations of results. With an increase of dynamics, the difference between the update frequency gets larger, due to an increasing lack of close information.

The differences SWF analysis shows between the densities is quite comparable to that made in the basic experiment.

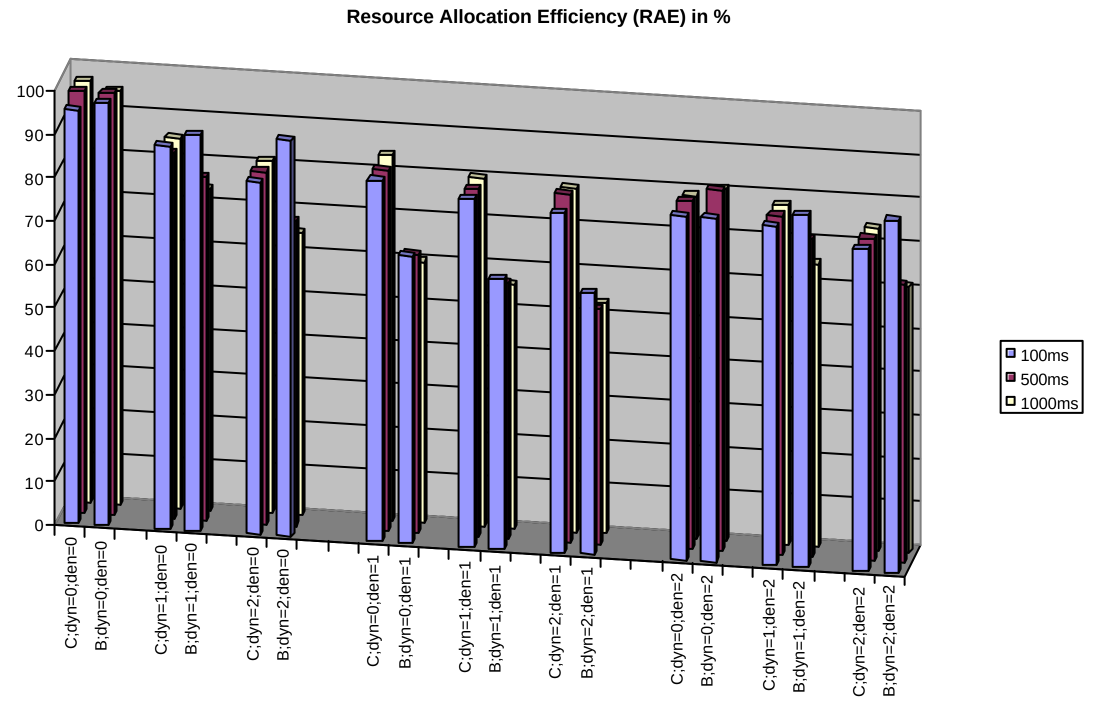

Figure 19: Effects of MSC Update Time on the RAE

Concerning the RAE, similar results were expected and achieved, which back the findings of the SWF. It can be seen that SWF and RAE are in a strong correlation and all observations made concerning SWF can be found in RAE as well.

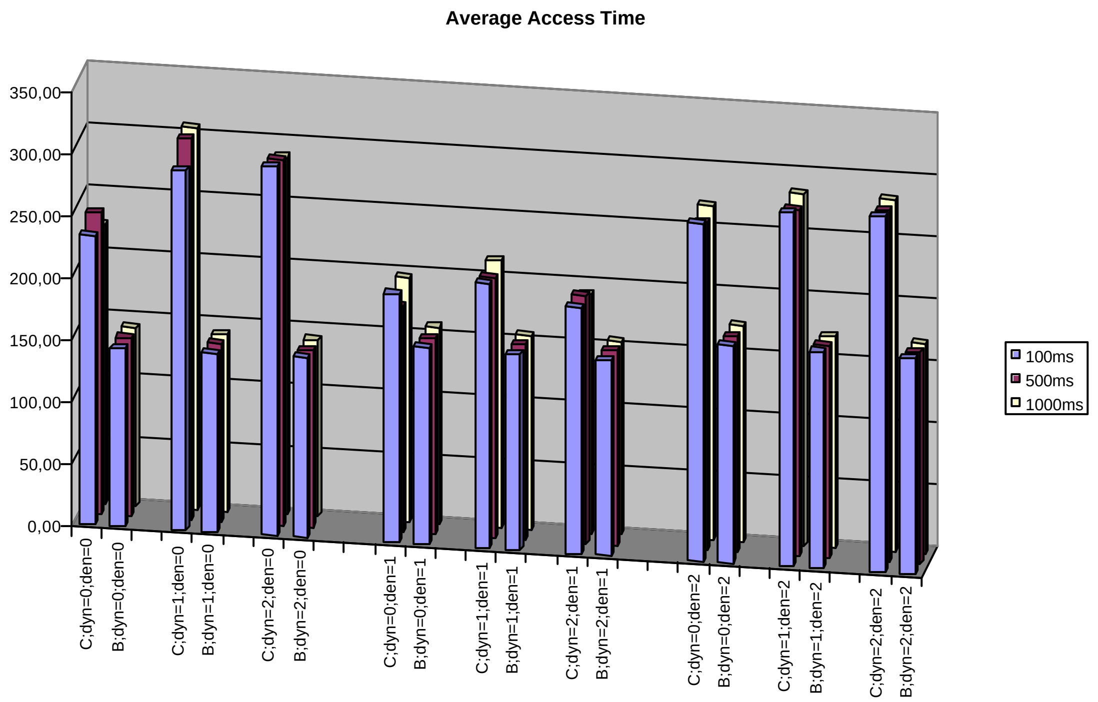

Figure 20: Effects of MSC Update Time on REST

The response times results of the different experiments showed no great changes; this parameter seems to be independent to the REST. This independence could be explained by the ability to easily select another SC to provide the service when noticing a node failure. Therefore, possibly update times do not affect REST.

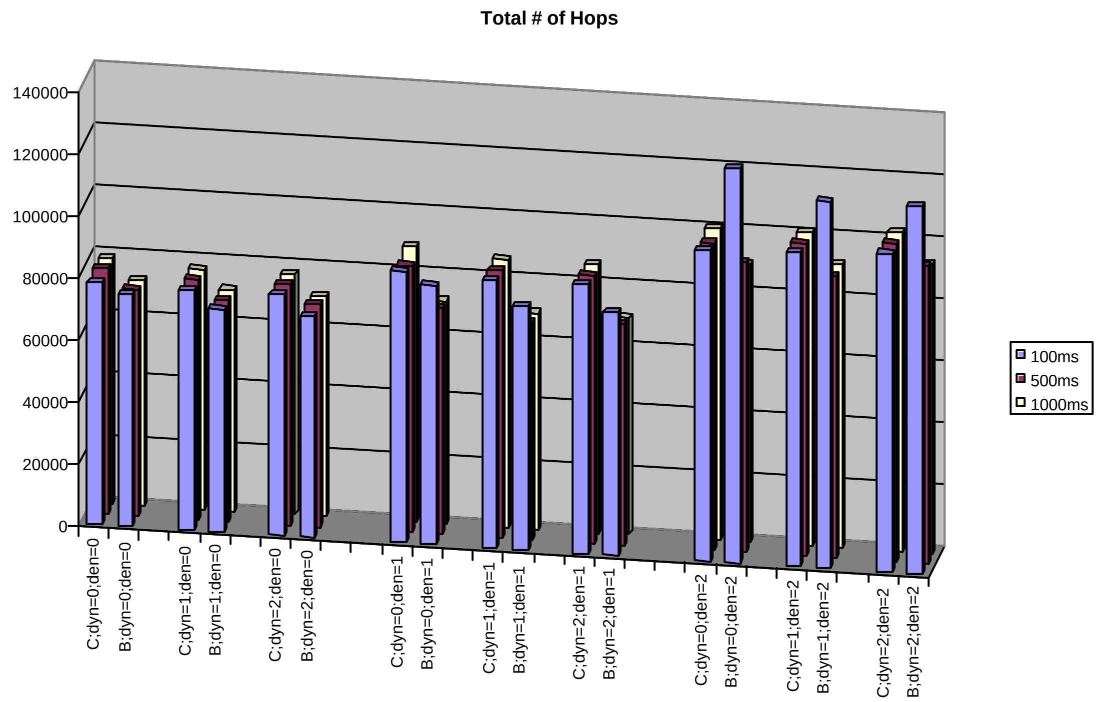

Figure 21: Effect of MSC Update time to hop count

The number of hops is affected by changing the MSC Update time. This has been expected, as a higher MSC Update time increases the number of messages sent in the network. Hence a decrease of CC can be marked when lowering the frequency of update mails (increasing update time).

Overall, the variation of the MSC Update time leads to some changes in the result of the Baseline model, comparing it with the basic experiments. However, the relation between Baseline and Catallactic results (e.g. when looking at SWF or RAE) does not substantially change. The inferior performance and the rejection of some of the hypotheses is thus not a function of the MSC update time.

3.3 Effects of changed MSC Allocation time

In another sensitivity analysis, we changed the allocation time of the Master Service Copy (MSC) from 125ms over 250ms and 500ms to 1000ms (default). The MSC Allocation Time describes the cycle time of the MSC, that is between the reception of the first request message and the computation result of the resource allocation. Any change in this parameter should only affect the outcome of the Baseline model, as the Catallactic model does not use the MSC. An expected outcome is that slower update times will lead to inconsistent information at the MSC and therefore to misallocations.

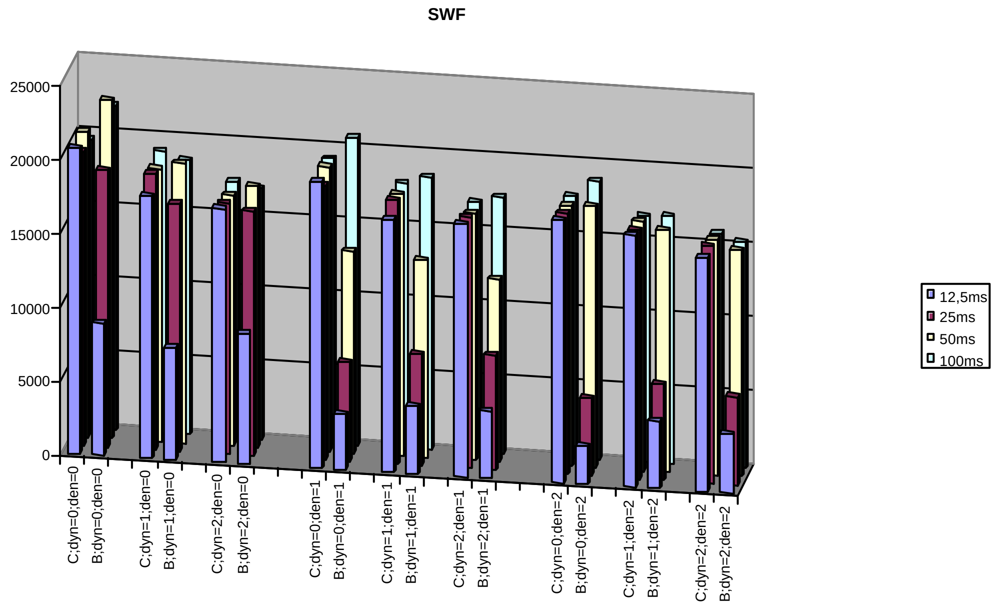

Figure 22. Effects of MSC Allocation Time on the SWF

Regarding the SWF, the MSC Allocation time has great influence on the SWF. This shows that the MSC needs a greater period of time to collect supply and demand and match it in an efficient way. It can be observed that social welfare decreases with dynamic and density in Baseline model and performs in the default case of 100ms similar to Catallactic behavior. In medium and high dynamic regimes a higher MSC Allocation time leads to lower SWF results. This shows that increasing dynamics leads to misallocations and thus to a lower ability of the application layer network to fulfill the individual goals of the participants.

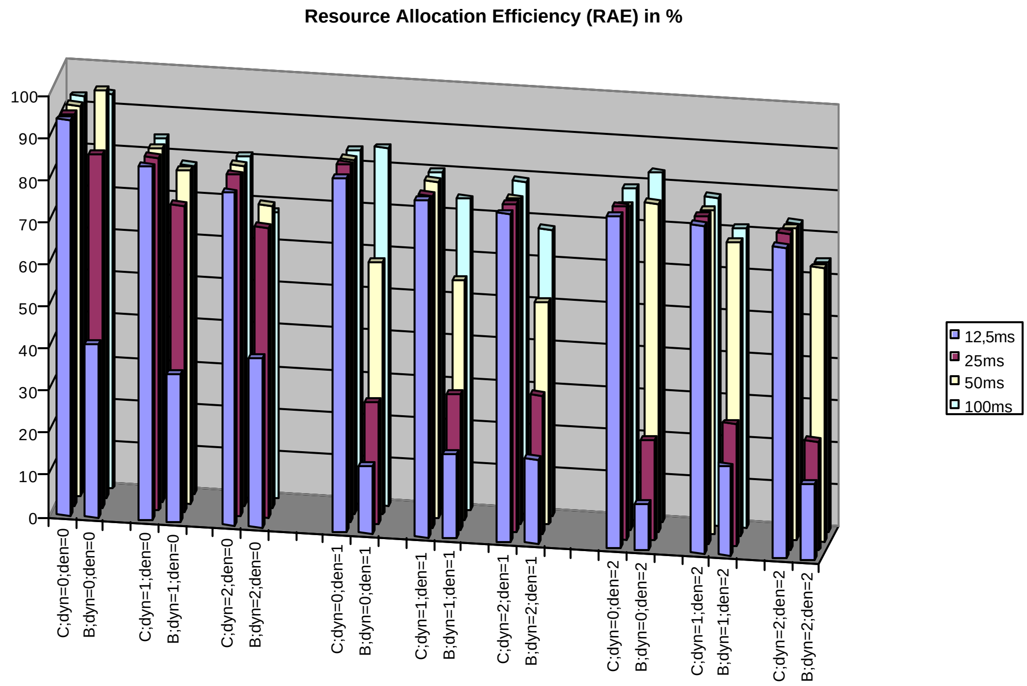

Figure 23. Effects of MSC Allocation Time on the RAE

Concerning the RAE, similar results were expected and achieved, which back the findings of the SWF. Results decrease with an increase in dynamics and density. As it can be seen, changing the allocation time affects RAE in the Baseline scenario. This might be a consequence of a certain time needed to get sufficient information to match supply and demand. RAE decreases rapidly in all scenarios when nodes often fail.

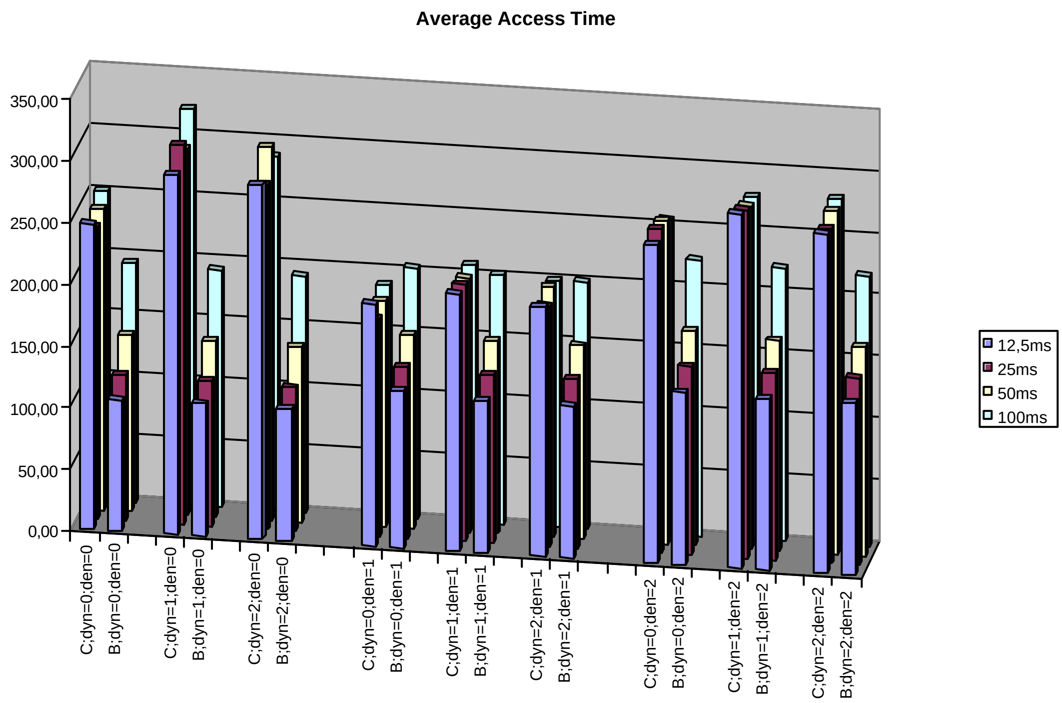

Figure 24. Effects of MSC Allocation Time on REST

Increasing the MSC allocation time directly means incrementing the REST. Naturally access times increase with a higher allocation time period as allocation time period adds to the service detection time. This can be observed over all set topology parameters and the difference is nearly independent from topology.

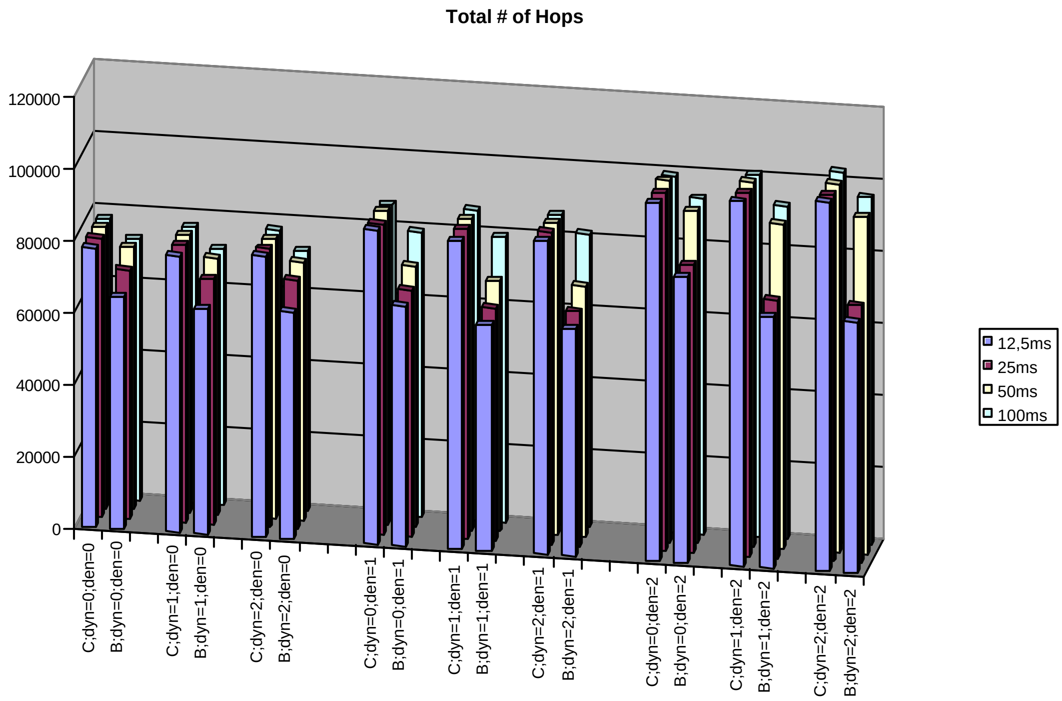

Figure 25. Effects of MSC Allocation Time on the number of hops (CC)

The number of hops is affected by changing the MSC Allocation time. This has not been expected, as the MSC Allocation time should influence the timing aspects of the network, but not the topology. It can be assumed that in combination with higher allocation times more accepts and money transfer mails have to be submitted to the bidding/asking nodes, which goes along with better RAE results.

Overall, the variation of the MSC Update time leads to some changes in the result of the Baseline model, comparing it with the basic experiments. However, the relation between Baseline and Catallactic results (e.g. when looking at SWF or RAE) does not substantially change.

3.4 Effects of allowing SC migration

The following graphics show the effects of an enabled migration to RAE, SWF, Access Times (REST) and Communication Costs.

Architectural SC are conceived to move their location. This could be done by bilateral negotiations like the service provision conversations. The SCs could e.g. move to decrease the distance between themselves and the demand sources and to overcome congestions to achieve a higher profit. Hence, the decision to move is made upon an economic calculus.

Migration is only implemented for Catallactic model, so only the left bars are considered.

To demonstrate the effect of the migration, a demand queue was modelled that changes its point of origin over time. So it can be mapped a change of the direction of the demand (oscillating demand), caused by time changes; e.g. during western European nights, the demand will go down and the main demand will come from the US, and vice versa). The first four graphics show the experimentation results with a static (non-moving demand queue).

The experiments are launched with different network parameters as before, so the results can only compared to the above experiments with caution, nevertheless the following graphics evidently show the effects of migration.

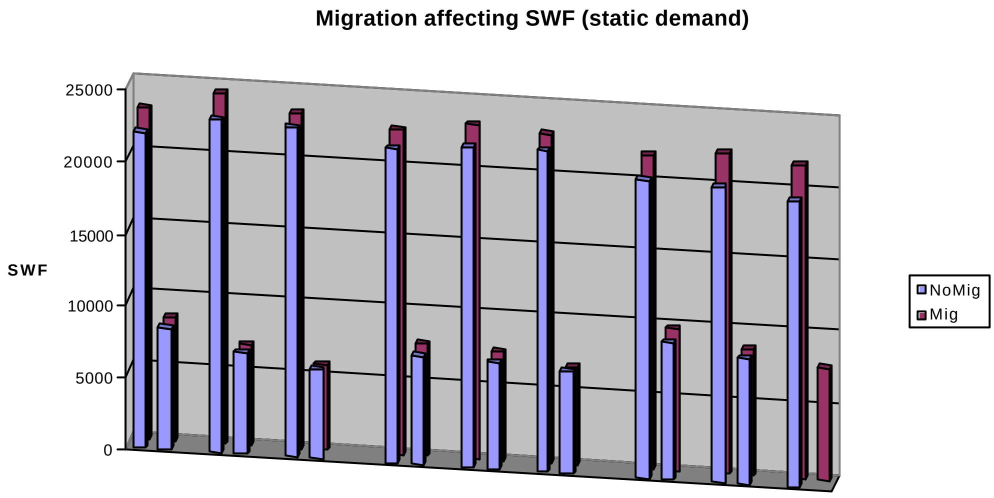

Figure 26: Migration affecting SWF (static demand)

The results show that migration has only little, but always positive impact on the SWF and adds a stable value to the SWF without migration. This could be an effect of an increase of found services within the predetermined hop counter as more SCs are in the range. Thus an increase of successful negotiations takes place. This should be the result of migration, SCs cope with congestions and try to overcome these by relocating.

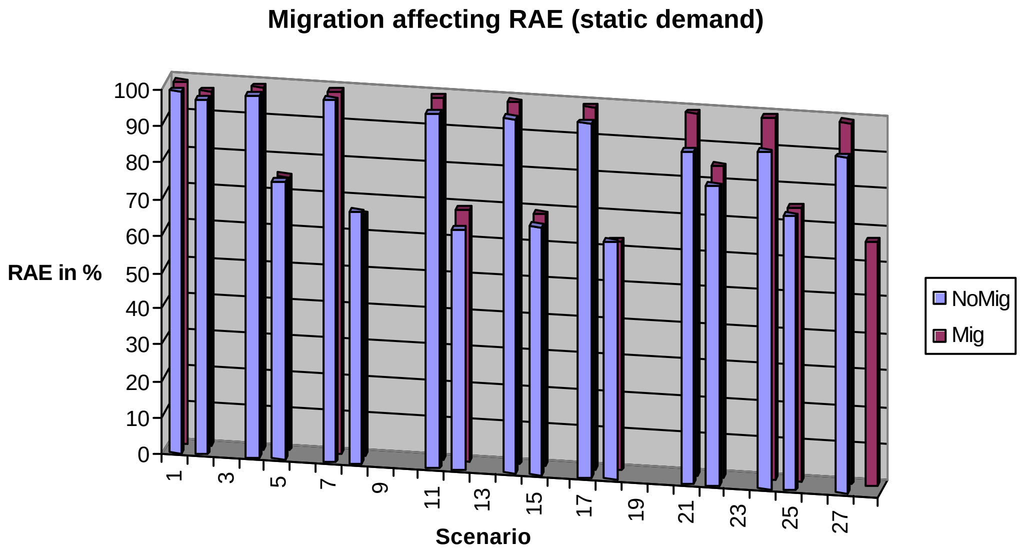

Figure 27: Migration affecting RAE (static demand)

The migration has an interesting impact on the RAE. It can compensate the losses of the greater distribution/higher density of the Service Copies. The argumentation here is very similar to that made in SWF. SCs overcome bottle necks (of resources or network connections) and relocate to earn more profit.

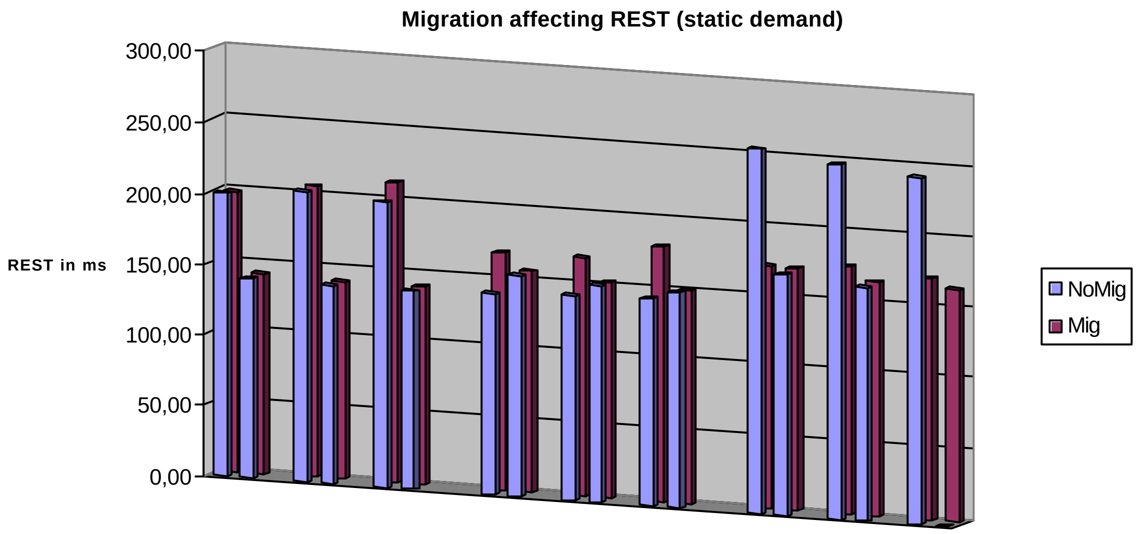

Figure 28: Migration affecting REST (static demand)

Low density and high density topologies can be transferred by each other. The migration compensates the effects of the distribution of SCs and achieves balanced REST results. It is assumed that migration will lead in all density scenarios to one equal density.

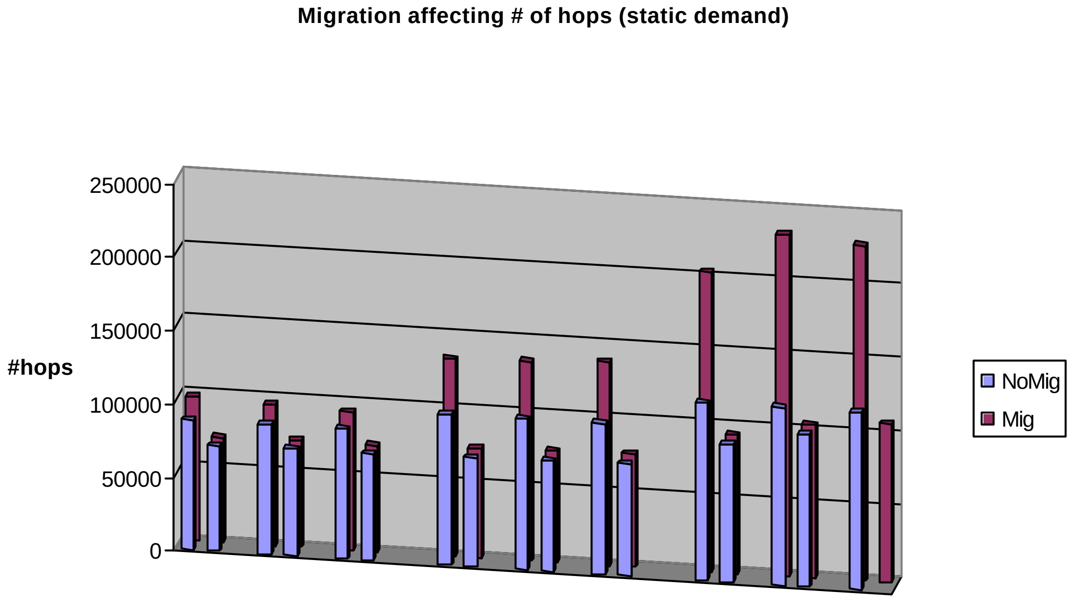

Figure 29: Migration affecting # of hops (static demand)

The number of hops is increasing in higher density networks in the migration scenario. This comes from the higher number of negotiations for storage that is needed for coordination of migration.

To point up the migration in a moving-demand scenario, the next four graphics show the results of an experiment with a dynamic demand. It should be observed that migration will perform even better with a moving demand as it can compensate those one-sided charges.

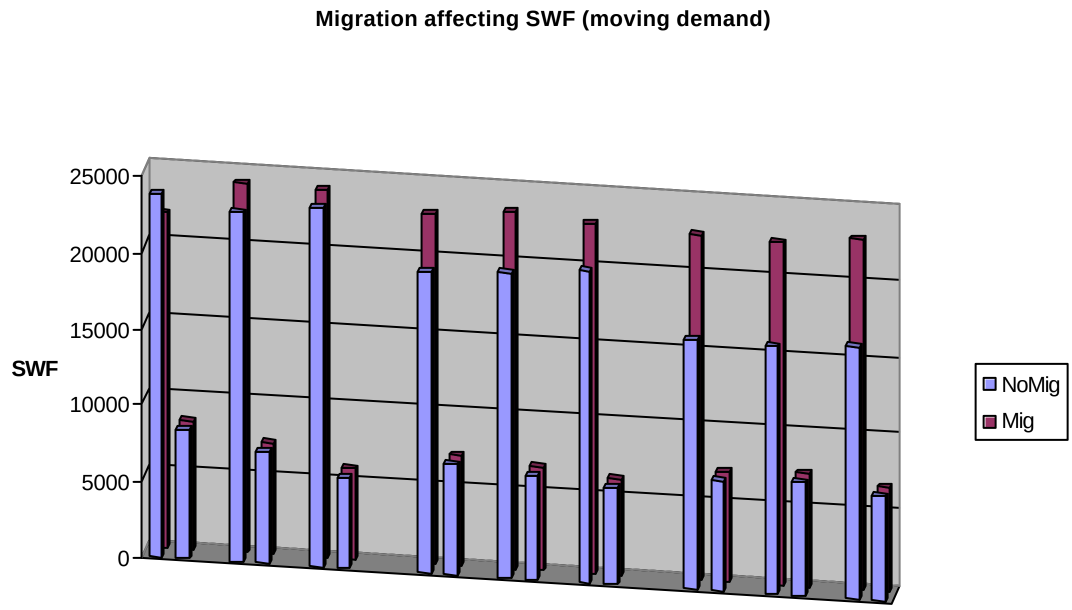

Figure 30: Migration affecting SWF (moving demand)

In comparison to the non-moving demand, the SWF is decreasing significantly without migration. This might be a result of changing demand sources and unreachable Service Copies. This effect is even strengthened by a moving demand. The migration can nearly compensate this behavior completely.

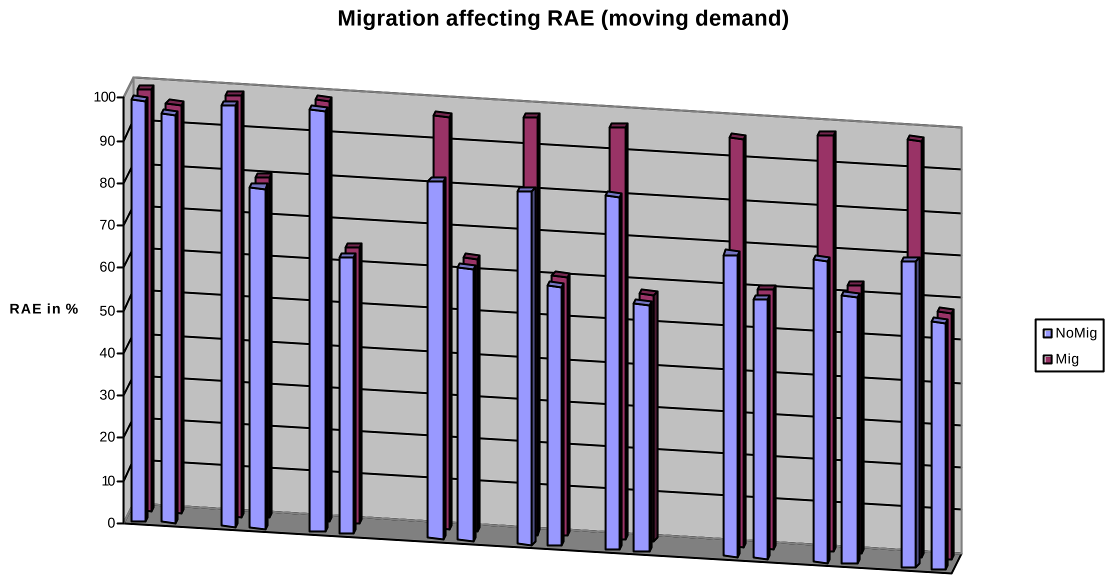

Figure 31: Migration affecting RAE (moving demand)

The results of measuring the RAE go along with those witnessed in the SWF contemplation. The argumentation is similar to those made in Figure 30.

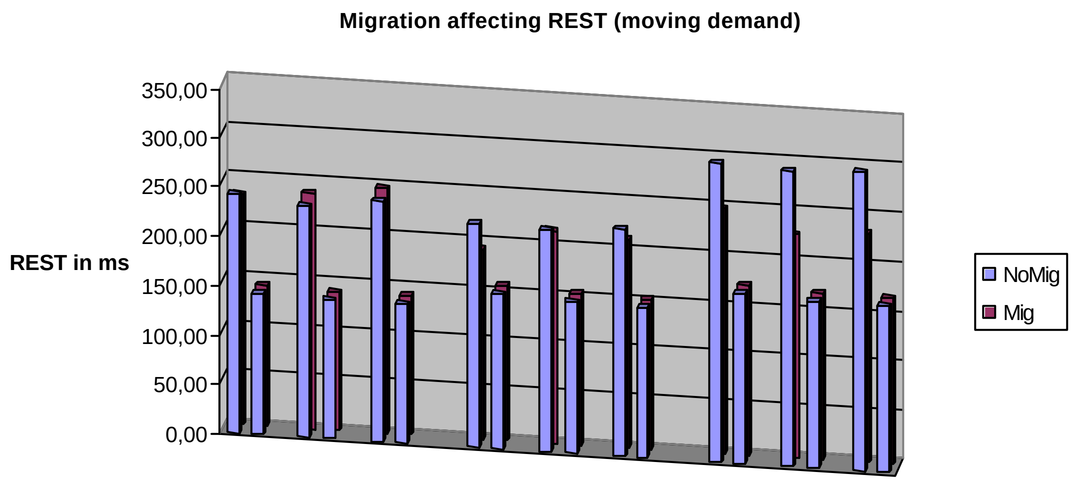

Figure 32: Migration affecting REST (moving demand)

In comparison to a static demand, REST increases. This behavior is reduced by migration for the explained reasons.

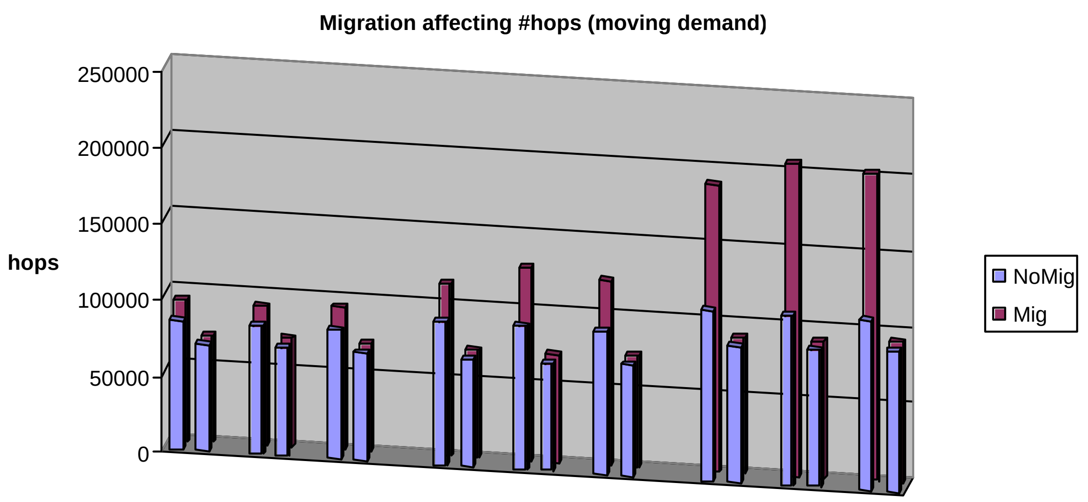

Figure 33: Migration affecting number of hops (moving demand)

In comparison to a non-moving demand the number of hops is higher. This could be the result of the increasing number of negotiations made for storage because of the changing demand sources.

It can be clearly seen, that migration has positive effects on the experimentation results. In the case of a moving demand RAE and SWF decrease extremely and migration can compensate this behaviour nearly completely.

4. Summary and Discussion

The following hypotheses have been described in the project proposal:

1. In a quasi-static system (low density, low dynamics) the Catallactic system operates providing nearly equal social welfare, with resource allocation efficiency slightly less than in the Baseline system, slightly more bandwidth utilization that in the Baseline system, and with a reaction time slightly longer than in the Baseline system.
2. In a very dynamic system (high density, high dynamics) the Catallactic system operates providing greater social welfare, caused by greater resource allocation efficiency compared to the Baseline system, with less communication cost, and with shorter reaction time.
3. In a low node density system with high dynamics, the Catallactic system operates at slightly less social welfare, caused by less resource allocation efficiency and slightly higher communication cost.
4. In a high node density system with low dynamics, the Catallactic system operates at equal or better social welfare, provided by a slightly less resource allocation efficiency than in the Baseline system, but even less communication cost than in the Baseline system.

The discussion presented in this section will relate the results described in section 2 and 3 to these hypotheses.

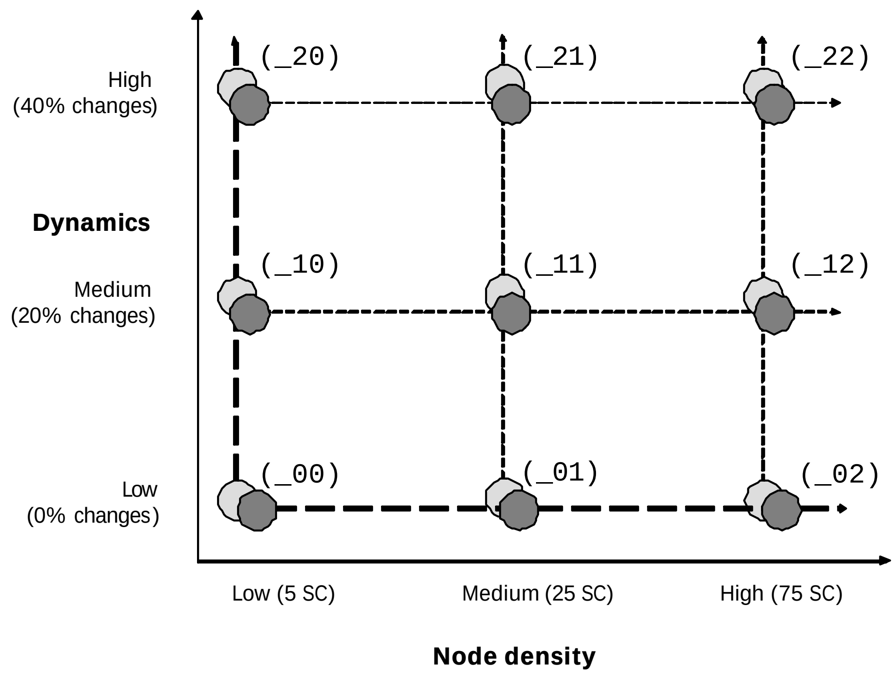

Figure 34: CATNET Experimental Setup

4.1.1 Social welfare utility

The main statements of the hypotheses described above with regard to the social welfare utility criterion (SWF) were, that the SWF value for Catallaxy is

- correlated with both increasing dynamics and increasing density, so that the main gradient of increasing SWF runs diagonally from the lower left to the upper right corner of Figure 34, but
- more correlated with density than with dynamics. In other words, an increase in density is more significant for changing the SWF value, as an increase in dynamics.

For the Baseline approach, the hypotheses have to be considered in reverse. The Baseline approach should perform better than Catallaxy in the case of decreasing dynamics and decreasing density. These findings have been found to be correct in general, but are affected by the concrete parameter setup and network topology. The current results thus present a mixed bag of confirmations and falsifications of hypotheses, and it is not possible to make a positive statement about the applicability of Catallactic resource allocation to application layer networks, on the ground of these experiments alone.

However, the SWF criteria achieves the intended goal of being able to assess the performance of both Baseline and Catallactic models, as they show notable differences.

Figure 35 shows the value of the SWF parameter in the different scenarios. The behavior of the Catallactic system changes rather smoothly from one scenario to the other, showing a straight decrease with increasing density and dynamics, while the gradient in Baseline system is higher and less regular. For Baseline, a decrease of dynamics leads to a decrease of SWF, while for density an unexpected extremum exists in the medium density regime. With the help of sensitivity experimentation, it could be shown that this effect can be influenced by the MSC Allocation Time. Concerning only the current parameter setting and network topology, the Baseline SWF seems to be correlated to the MSC Allocation Times as shown in the respective annex chapter.

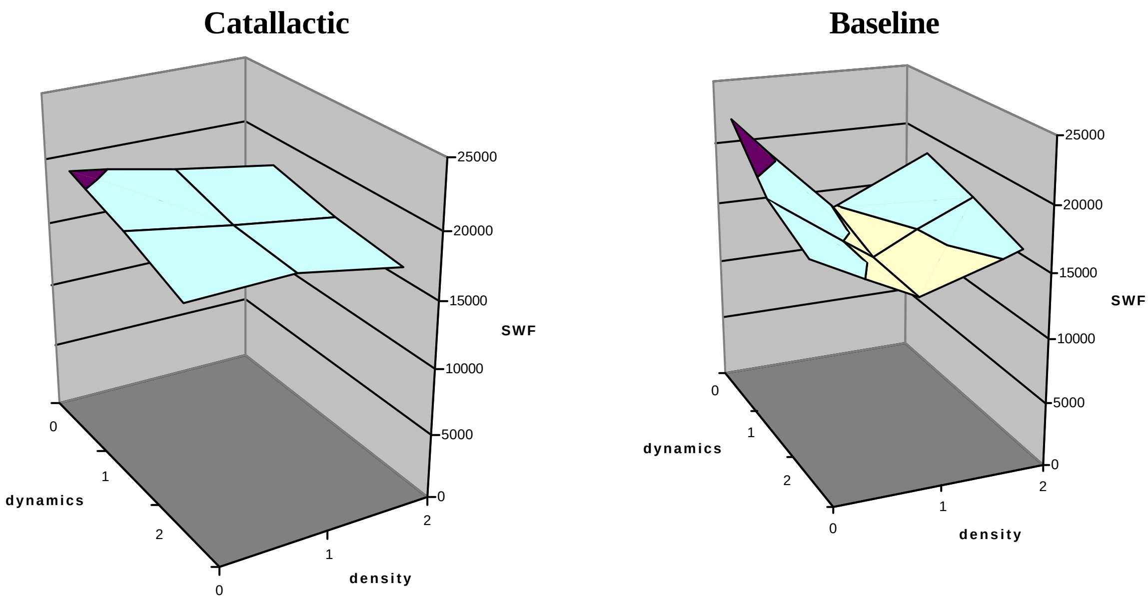

Figure 35. SWF of the Catallatic and Baseline in the 9 scenarios.

Overall, the explanatory power of the SWF criterion is theoretically sound, but depends practically on (too) many factors. From an experimental viewpoint, the main advantage of using SWF for measuring the ALN performance is that it incorporates all costs and revenues the agents incur by participating. An increase in SWF indicates that *overall* the utility gain of all agents has increased. For interpretation, this feature is also the biggest drawback, as too many interdependent factors hinder a detailed cause study and explanation of the phenomena found (see section 4.1.5 for a discussion on the interdependencies).

4.1.2 Resource Allocation Efficiency (RAE)

The hypotheses for resource allocation efficiency (RAE) expect a performance similar to the SWF criterion, but not equal. In particular, only the high dynamics/high density scenario is expected to show a higher RAE for Catallaxy than for Baseline. Both the low dynamics/low density and the high dynamics/low density scenario are expected to show a slightly less RAE, while the low dynamics/high density scenario should show a significantly lower RAE for Catallaxy than for Baseline.

Again, the outcome presents a mixed bag of results with confirmation and rejection of hypotheses, very much aligned to the development of SWF (see section 4.1.5 for an explanation). In the quasi-static scenario (_00), both Baseline and Catallactic RAE were close together, with the Catallactic RAE a little bit higher. The hypotheses H1.2 could thus not be confirmed.

For the highly dynamic scenario (_22) the difference between the superior Catallactic RAE value and the Baseline value was just a little bigger. Despite the possible confirmation of hypotheses H2.2, the difference was not considered big enough and the hypotheses could thus not be confirmed.

For the low node density scenario (_20), surprisingly the Catallactic RAE value is even higher than for the highly dynamic scenario. Hypothesis H3.2 is thus rejected, but the reason becomes not clear enough to make a final statement.

The high node density scenario (_02) shows a RAE value which is comparable to the highly dynamic scenario. However, it is not clear whether this result correlates with high density.

In Figure 36 the value of the RAE parameter in all considered scenarios is shown. It can be seen that the RAE of the Catallactic system changes smoothly from one scenario to the other. The RAE achieved in the high density scenarios is almost independent of the dynamics, but shows some negative correlation to the factor density. The RAE of Baseline is less regular. Higher dynamics decrease the RAE achieved. Observing density, the same behavior is exhibited like in SWF. Obviously a minimum of RAE can be noticed in the medium density scenarios which is suspected to be due to the same explanation.

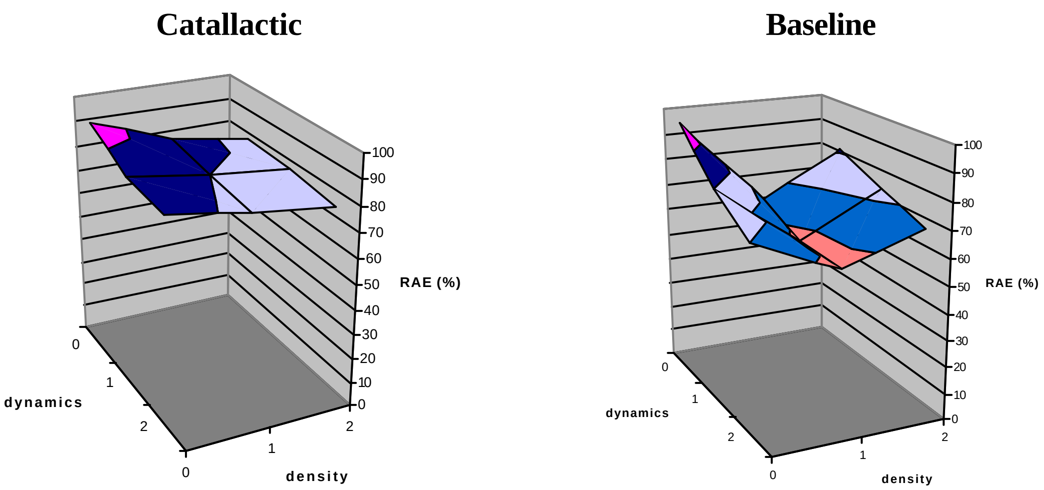

Figure 36. RAE of the Catallactic and Baseline in the 9 scenarios.

4.1.3 Response Time

The main statement of the hypotheses with regard to the response time (REST) was, that the REST value for Catallaxy is positively correlated with the combination of increasing dynamics and increasing density. There have been no statements made for the isolated change in only one dimension.

In the quasi-static scenario (_00), REST was much longer in Catallactic than in Baseline, therefore hypotheses 1.3 could not be verified. This proves the superiority of the Baseline approach in a stable non-changing environment.

The high dynamic scenario (_22) presents an average longer REST in comparison to Baseline, which was not proposed by the hypotheses. Like discussed before, it is considered that due to a higher probability of node failures more negotiations have to be initiated to establish successful service provision.

For scenario _02 (low node density) the results showed higher REST in Catallactic approach. This could be due to high service detection endeavours.

For the high node density scenario (_20) it was expected to obtain longer Response Times in Catallactic approach. This could be verified. REST was about 95% higher. This could be the effect of a high supply of services so that the client could make several negotiations to get the best offer.

Figure 37 displays the value of the response time (REST) in the different scenarios. The response time of the Catallactic system is influenced both by the dynamics and density of the simulated scenario. In this model, a minimum is shaped at medium density and a maximum at medium dynamics.

In Baseline, the response time grows with the increase of node density and shows independence to dynamics.

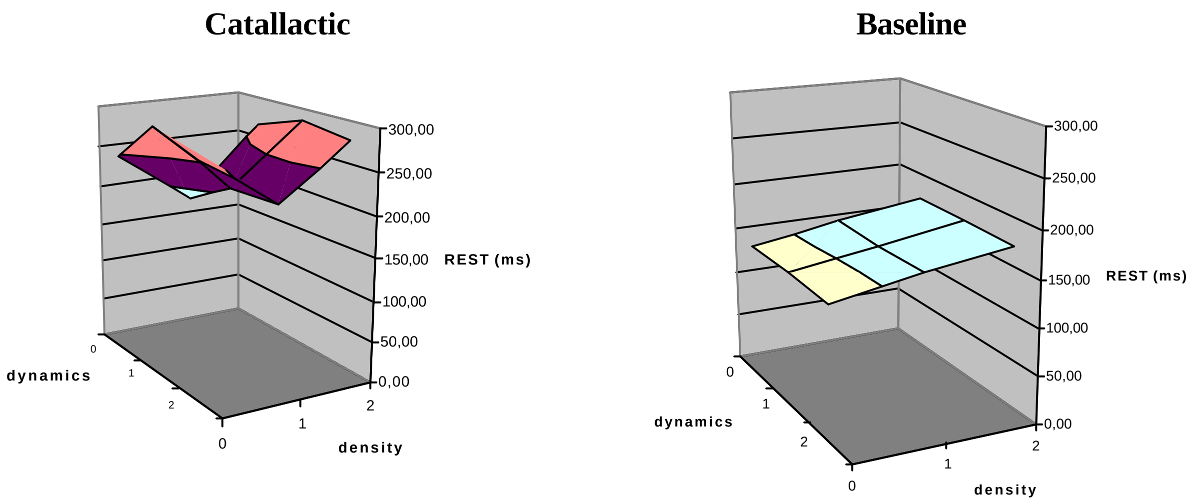

Figure 37. REST of the Catallatic and Baseline in the 9 scenarios.

To put it in a nutshell, the Catallactic REST in all scenarios was outperformed by the Baseline approach, sometimes significantly. This is an unforeseen result. It is supposed to be the effect of time-consuming negotiations in high density scenarios and lasting decentralized service discovery times in low density scenarios. In addition to that, in high dynamics the high number of conversations necessary for service provision will cause longer response times.

4.1.4 Communication Cost

The main statements of the hypotheses with regard to the communication cost (CC) were, that the superiority of CC value performance for Catallaxy is mostly correlated with increased node density. Increasing dynamics alone, as in the low density/high dynamics scenario, leads to higher CC, while increasing density alone, as in the low dynamics/high density scenario, leads to lower CC. The combination of changes in both dimensions should lead to an over-compensation, where the effects of density override those of dynamics.

Resuming the results, in the quasi-static scenario (_00) CC were expected to be slightly more than in Baseline, due to its supremacy in static environments. This expectation could be verified.

The high dynamic scenario (_22) showed that CC are slightly higher in Catallactic. As discussed before, the non-static topology causes a higher, unexpected conversation emergence. Therefore, the hypotheses had to be rejected.

For low node scenario (_02) slightly higher CC were awaited. This could be confirmed. Considering REST it can be assumed that the difference between REST and CC is caused by long queues at the nodes because of congestions and not by long negotiations.

For the high node scenario (_20) communication costs are even higher, which falsifies the hypothesis.

In Figure 38 the value of the CC parameter in the different scenarios is presented in comparison to dynamics and density. It can be pointed out that the CC both of the Catallactic and the Baseline system increases with higher node density but does not change by altering dynamics. Catallactic shows a steady increase, whereas in Baseline model, an increase firstly comes up when changing from medium density to high density.

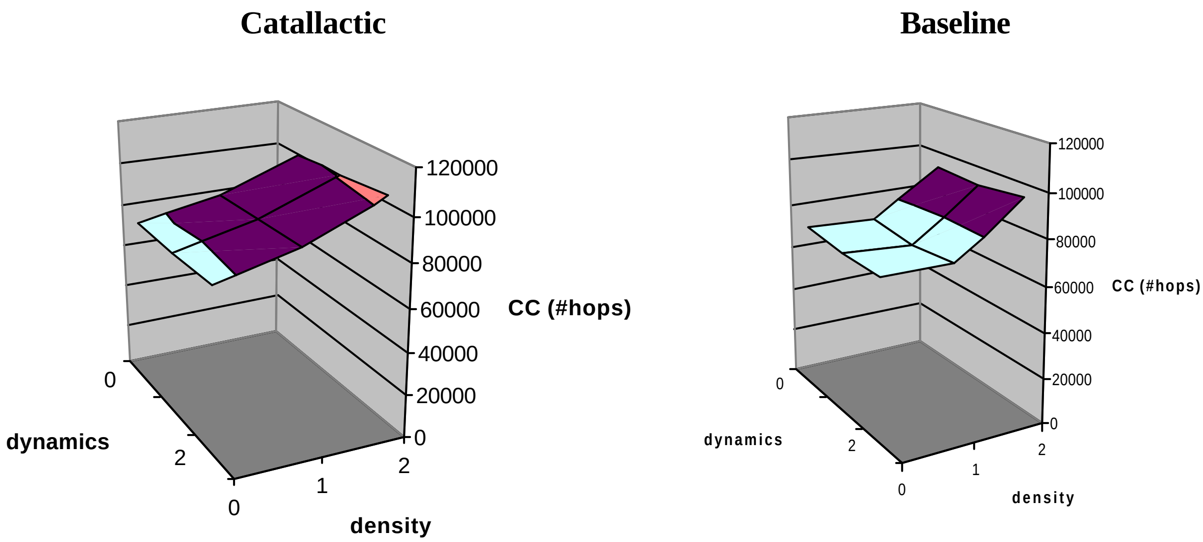

Figure 38. CC of the Catallactic and Baseline in the 9 scenarios.

Summing up, communication costs are higher in all scenarios in the Catallactic approach. This might be because of a higher coordination endeavour in the decentralized case. Interestingly this behaviour often leads to better results which have to be "paid" by higher communication activities.

4.1.5 Interdependence of criteria

At the start the project has assumed the social welfare utility (SWF), resource allocation efficiency (RAE), access time (REST) and communication cost (CC) criterion as being independent of each other and depending only on the variation of node density and node dynamics. The findings presented in this report, visualized by the respective figures, show that the criteria results are affected (1) by each others development and

(2) by other parameters than density and dynamics. This section investigates into the interdependencies.

In general, we have found that the Catallactic model becomes relatively superior over the Baseline model with increasing dynamics and increasing density, as indicated by the social welfare utility (SWF) criteria. But Catallaxy achieves this result at the expense of higher communication cost and inferior response times. The RAE criterion mirrors quite closely the development of SWF. This leads to the assumption that SWF is positively correlated with RAE, and inversely correlated to both CC and REST.

The interdependence of SWF and RAE is not such surprising, as both criteria are based on the ratio of successful transactions within a given time span. In the current implementation, Client agents and Service Copy agents both agree to a transaction only if the utility gain is positive. To this respect, every successful transaction adds to the SWF value; the more transactions are successful within the experiment in total, the higher is the total SWF at the end. This interdependence could be deliberately broken, if the negotiation strategy of the agents would allow them to conduct transactions with negative utility gains.

The interdependence of SWF and CC is not surprising in so far as the SWF value for each agent should already compensate all the costs and revenues incurred in the agent. To this respect, an increase in the communication costs should lead to a decrease in SWF, as for a given revenue situation the costs rise. This shows clearly in the figures shown in the previous section.

For the correlation between SWF and REST, the link lies in the number of transactions conducted in a given time span. A higher response time means also slower reactions on whether to take up a negotiation or not, which in turn affects the possible number of transactions during the experiment. As has been indicated above, the number of transactions directly and positively influences the SWF.

In total, the dependencies between the criteria came out clearer than expected. Despite the sometimes tautological findings on the experiments, these dependencies are an *emergent feature* of the simulation, as no direct relations between them have been programmed. To this respect, we regard these correlations as positive insofar as they support each other and give the statements more consistency and weight.

At last, a note on the particular result of the medium density regimes. The figures presented in this report show clearly that in the results for all criteria there is a strong dependency to the value of density, with the best results for Catallactic with density=1 with any value of dynamics: the *x1* experiments. It seems that by changing some parameters other than node density and node dynamics, these (local) maxima or minima can deliberately be made to appear or not. In the current basic setup as shown here, it looks like a local optimum for network density or dynamics would exist. However, this may not be the case for other setup combinations, and thus can not be generalized to practical Grid applications without further research.

5. Conclusions and Outlook

The achievements of the project are:

1. the development of an experimental simulator for investigating different allocation mechanisms, allowing their comparative evaluation;
2. the definition of a two-dimensional problem space, defined by dynamics and density, which can be used to categorize application layer networks;
3. the conduction of an experimental evaluation study of two allocation mechanisms.

5.1 Regarding the development of an experimental simulator

There are several simulators for Grid and Grid-like networks available [1]. In most of these, the implementation of the networking, the application and the allocation mechanism build a proprietary whole. The evaluation of the performance of different allocation mechanisms is thus not possible. The CATNET simulator has been designed from the beginning to provide abstract interfaces for messaging and allocation computation, so that both centralized and decentralized allocation mechanisms can be used.

For the decentralized allocation mechanisms like in the Catallactic scenario, service clients and service providers communicate using a well defined messaging protocol. The decision-making is encapsulated in the client and provider classes using a heuristic-adaptive strategy, which can easily be modified or replaced without side effects to other classes in the system. It is thus possible to try out other rule-based, heuristic-adaptive or even game-theoretic strategies.

For the centralized allocation mechanisms like in the Baseline scenario, the modification is even easier, as it only effects the internal processing of the MSC class. It would be possible, for example, to experiment with a heuristic “lowest price choice” approach or even a stock market/equilibrium price approach, without having to change any other object in the simulation.

5.2 Regarding the categorization of application layer networks

To evaluate the project findings and to map them to realistic scenarios, we have categorized the different existing projects, e.g. under the labels of Data Grids, Computational Grids, Peer-to-Peer Networks, using two abstract dimensions, node density and node dynamics.

The node dynamics dimension looks at the frequency of changes in the application layer network, characterized by the appearance and disappearance of service offers and demands. This is different from the network layer, where network nodes or communication links break. An example of high node dynamics is the high frequency of nodes connecting and disconnecting from the P2P network; this connecting act is (mostly) a deliberate act of the participants, starting up or shutting down the P2P client.

For low node dynamics, resource nodes are considered to be highly available and permanently connected to the fixed network. This would be a feature of top-down designed networks like DATAGRID, or in an extreme case of parallel computation networks. In between low and high node dynamics, it is possible to find numerous examples, which can be simulated using CATNET.

The other categorization dimension is density, which reflects the average distance between service providers and clients. Again, this measure is on the application level rather than on the network level. It does not measure the connections between particular network nodes, but abstracts from the particular network topology, although this still has a measurable effect on the outcome.

In summary, the two dimensions dynamics and density have been found to describe different service networks and explain their difference on an application layer level, as opposed to pure communications network measures.

It is subject to future research (e.g. the theoretical approach, see below in the outlook chapter) to state whether these two dimensions are sufficient to explain, or whether additional dimensions have to be taken into account. These findings could in turn lead to better simulation setups, and also for planning and deploying Grid networks.

5.3 Regarding the conduction of an evaluation study

The outcome presents a mixed picture which needs further exploration, to overcome artefacts and limitations of the simulation software. The current results are valid for the particular setup parameter values, the specific topology and the measurement procedures.

The absolute value of the criteria performance is highly sensitive to that particular setup, and can easily be changed by a variation in the input parameters. The absolute value can thus not be used to provide performance statements. However, the relative value and rations between Catallactic and Baseline results have been found to show similar patterns, which we are able to repeat using the same simulator settings.

The results presented in the figures in this report are singular values from on representative experimentation run. For a statistically solid conclusion to further support our findings, the values of several experiment runs would have to be aggregated, and the average and variance of the values presented. For technical reasons, these aggregations were not possible to obtain during the runtime of the project. However, as the results values are repeatable and very similar for each experiment, the findings presented in this report are regarded as representative.

Despite these restrictions, the information provided by relatively comparing Baseline and Catallaxy within the same parameter setting, the same topology and the same measurement procedures comes to clear, repeatable conclusions about their respective performance.

One of these repeatable artefacts is the apparent “relatively optimal” performance of the “medium” regimes, especially for medium density, in all experiments. In contrast to our expectations and to the hypotheses which aimed at the “corners” of our problem space, it seems here that an optimal state of the network exists where both dynamics and density are somewhere “of average value”. This may lead to an engineering research question, which asks if an optimal state of the network exists, how can the dynamics and density of the network adapt to reach that state?

For the dynamic part, that answer has to be given by the technical implementation level: for a "naturally" low dynamic environment, it may be positive to deliberately add dynamics and faults so that the adaptation process in the network nodes can continuously improve; for a naturally high dynamic environment, some persistency and caching mechanisms might be implemented that implement self-healing and fault-tolerant capabilities.

For the density part, the answer lies in the application layer rather than the technical implementation. As the density can be changed either by deployment of new Service Copies or by migration of the existing SCs, the concrete implementation is probably dependent on the application properties. For CATNET we conducted a sensitivity analysis on using the migration capabilities or not. This experiment showed an increased performance for the SWF and RAE criteria, if SC migration was allowed. In this case, the change of the density towards a perceived optimum is an "emergent" feature, meaning that it is a whole-system phenomenon, achieved by deliberate action of the single network participants who work for their own good without taking possible effects on the system scale into account. This finding opens the door to the scientific field of "complex adaptive systems", which is pursued in the European Union's 6th Framework programme as a Future and Emergent Technologies Proactive Initiative.

Although the main simulation parameters which we vary are the node dynamics and node density, we observed that actually the design space which could be considered for both systems is much larger. Some indication of how the adaptation of other parameters affects the simulation is shown in the "sensitivity experiments" as shown in the respective chapter for demand frequency, migration, MSC allocation time and MSC update time. Other parameters we have considered to evaluate are, for instance, the effect of scale on the coordination mechanisms, the influence of particular characteristics of the demand queue, design parameters of the Baseline system to handle highly dynamic environments, and parameters of the strategy used in the Catallactic coordination to determine prices.

5.4 Outlook

For the future pursuit of the research questions of the CATNET project, two main directions can be followed to come around the obvious dependency of numerous setup and simulation variables.

In principle, one approach would be to decrease the number of variables until the outcome is provably dependent only on the dimensions; that is, the measurement criteria can be theoretically shown to be function values only of node dynamics and node density. This can be called the "theoretical approach".

The other approach would be to choose values which are realistic and lead to results comparable to those real-world implementations. The numbers of variables, the parameter setup and the network topology have to be taken from a real-world implementation, with the goal of making the Baseline scenario indistinguishable from that implementation as a reference. This can be called the "practical approach".

For the **practical approach**, the most effective way of building such a technical simulation would be to create resource allocation middleware, that also allows taking measurements, and to implement it e.g. in a GLOBUS-based real project under "live" conditions. The "technical simulation" of the Grid network would then be the Grid network itself; both the parameters and the parameter settings are taken from the real world application and the current implementation of the GLOBUS network. The number of variables which can be changed in the experiment decreases greatly and allows concentrating on further design aspects of the middleware. The criteria and measurement would also be derived from the goals and constraints of the Grid application, so that a "superior" criterion can be assumed to exist.

The downside of the practical approach is that an "evaluation" with a theoretical claim is nearly impossible to pursue. The middleware created is necessarily proprietary to the special technical implementation as well as to the application environment. The theoretical question whether a Catallactic, self-organizational allocation aspect is better for Application Layer Networks *in general*, can thus not be answered.

For the **theoretical approach**, the number of influential variables has to be decreased significantly in order to arrive at generic results. In the current CATNET implementation, the sensitivity of the results to the adaptation of variables like MSC update time or demand frequency is quite high. The results can change so much, that the judgment about confirmation or rejection of the simulation hypotheses is affected. For a theoretical statement about the performance of Catallaxy *in general*, these side effects need to be eliminated by reducing the number setup parameters. The goal is that only a change of node density or node dynamics should lead to significant result adaptation, while the variation of other parameters at least does not change the main statement regarding the hypotheses.

This theoretical approach could be realized using less functionality both in the network simulator and in the simulated application layer network, or by building the simulation on top of a known general-purpose simulation framework like Mathematica®. The downside of this approach is that the degree of abstraction is much higher than in the current CATNET project, so that a sound transfer of the statements to a realistic application is limited, even impossible. However, such a project would be able to give fundamental insight to a theoretical research field on the interface between computer networks and economics; despite missing practical significance, the possible scientific gain is regarded as quite high.

Some personal experiences of the CATNET project members finish this project report. During the runtime of the project, the development of the simulator, the design of the problem space and the conduction of the experiments have influenced each other more than we expected at the start. The joint writing and reading of academic papers and visiting academic conferences brought new knowledge into the project, which also led to modifications in these 12 months. All in all, the research topic, as precise as it was formulated at the beginning, began to move, the deeper we delved into the implementation of the CATNET Java classes using the JavaSim simulation framework.

Another obstacle we encountered was the sensitivity of the experimental results from other parameters than expected. This makes scientific reasoning complicated, but on the other hand opens doors to related research questions and provides associations which may be important for future Grid applications as well.

Under these circumstances, we regard the CATNET project as highly successful and as a good basis to pursue further research from here, either from a theoretical, practical or joint perspective.

6. Annex: Development of the parameters during experimentation

The development of the evaluation parameters during experimentation performs differently in Catallactic and Baseline approach due to the changed conditions of the models. The graphics below show this behaviour over time in a Catallactic and a Baseline experiment.

The price is represented in green, RAE in blue and Access Time (REST) is depicted red. The X-Axis illustrates experimentation time in milliseconds, the Y-Axis the

a) contract prices for two available services (multiplied by 6),

b) the RAE (multiplied by 5), and

c) the access time in milliseconds.

The price adaptation and the RAE values have been modified to fit into the displayed graph, while the results are not affected.

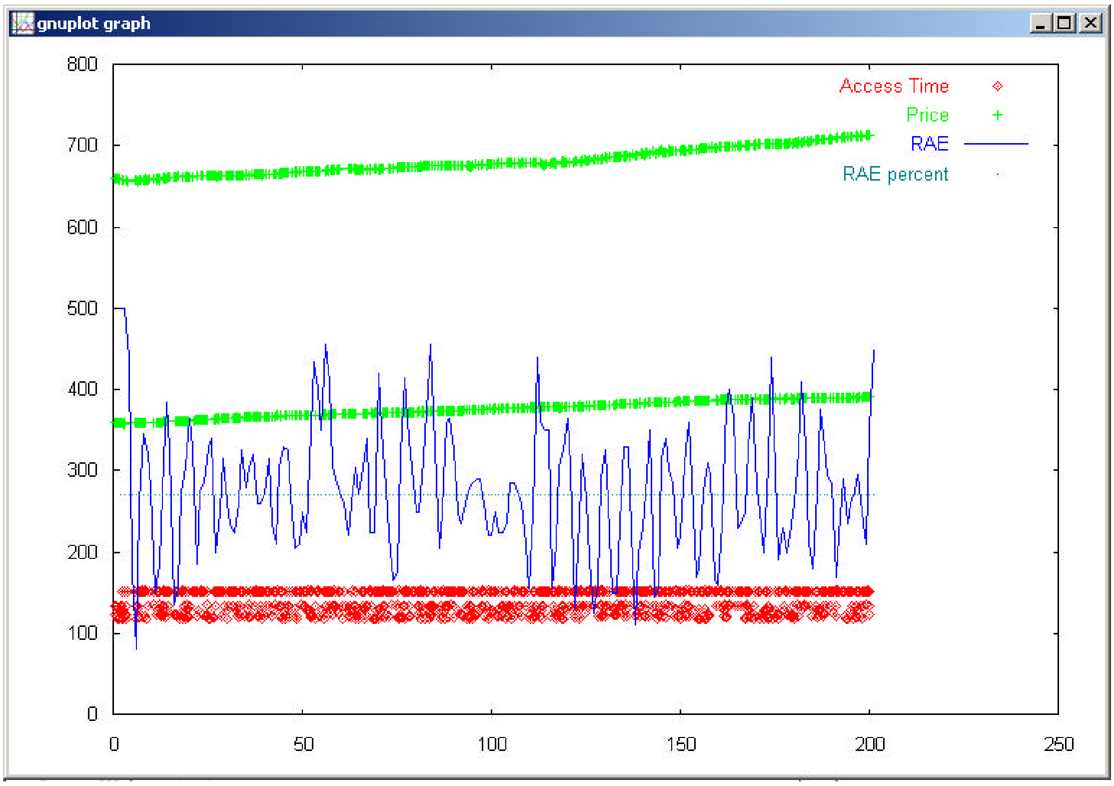

Figure 39: Baseline development of parameters during experimentation

For the Baseline model one can see in Figure 39 that prices do change a lot over time. There is a slight variation in prices. This may be subject to a successful price adaptation by the sellers; every seller agent increments its price perception over time and vice versa, while the clients always pay the cheapest prices, but have an inferior strategy to hold prices down.

RAE is in average about 50%. The high variance comes from the allocation cycles of the MSC.

The access time (REST) shows three different levels at about 140ms. Figure 39 shows a high density scenario, so SCs are very close to the clients and the three levels of access times come from the fixed network topology shown in Figure 3.

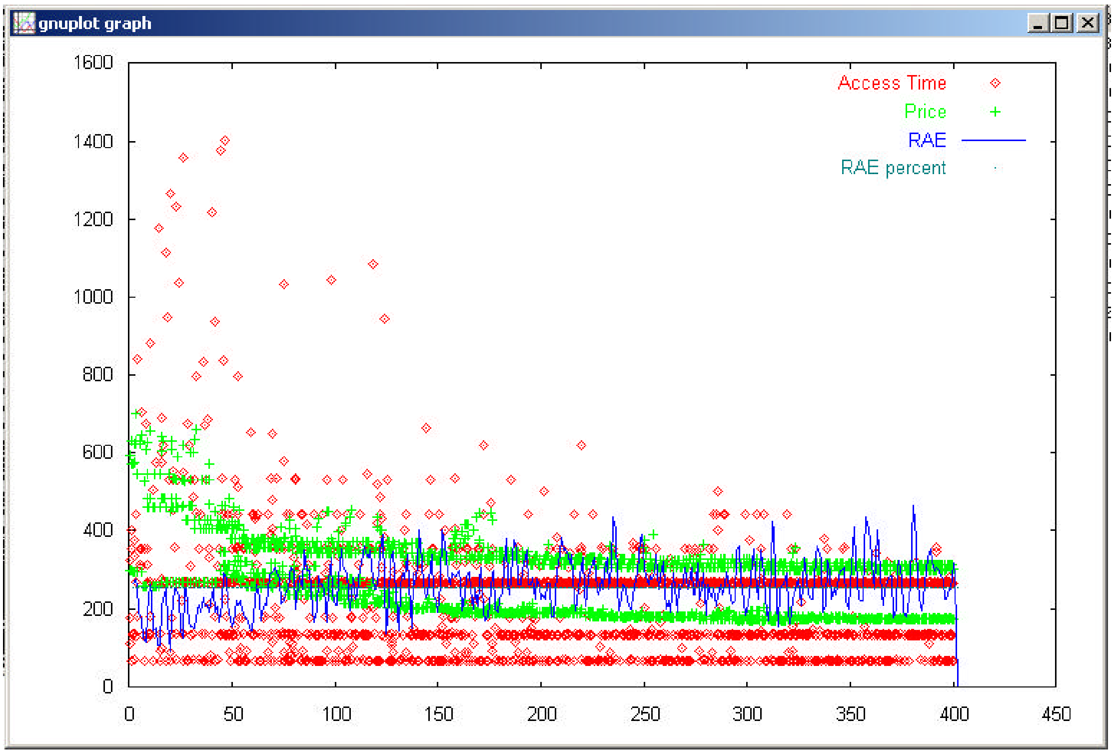

Figure 40: Criteria development during Catallactic experiment

In the Catallactic model contract prices are a result of continuous negotiations. Initially prices are set in the start script in a certain spectrum. Depending on a specific value in the strategy procedure, proposals are generated randomly, based on the price information obtained by the start script.

Initial bargaining prices adapt to the development of demand and supply. It was assumed that by using this simple heuristic strategy, an equilibrium price can be achieved for each item.

Agents adapt their price perceptions to generate a maximum profit. One can see, that a common level of (equilibrium) prices are established after 50ms, which is equal to approximately 250 demand requests.

RAE is right from the start extremely low. This might be due to a lot of failed negotiations committed because of too high price perceptions. After establishing the market (equilibrium) prices, it can be seen, that RAE climbs to a constant value.

The access time (REST) behaves similar. At first very high values appear. Partly access times of 1400ms emerge. When reaching the equilibrium prices, REST decreases to three lower constant levels. The hypothesis behind these three steady values is that in case of equilibrium prices less negotiation occurs and the REST represents the time distance to the three closest Service Copies (according to the topology).

In comparison to Baseline it can easily be discovered that the prices in Catallactic decrease. The Catallactic model represents a market, Baseline is more an allocation method; prices do change in a very slow form depending on a recent reject or accept.

There are no negotiations and therefore prices cannot change that rapidly like they do in the Catallactic model.

RAE appears in Baseline in a wider range, which might be due to longer periods without allocations made. If there is a central allocation issued by the MSC, RAE raises to a maximum value. If temporarily no allocation is issued, RAE will amount to 0%.

In Catallactic experiments, contracts and allocations are made by local decision-making, decentralized and continuously. Because of the higher number of negotiations, the variance is lower.

When comparing the REST it is obvious that access time in Baseline immediately positions at its average level. Due to the absence of costly price adaptations and negotiations access to a service can be achieved from the first request to the last.

To sum up, Baseline and Catallactic parameters perform totally different in its development. Baseline has no adaptation process and can establish better results in the beginning phase. Catallactic can catch up when establishing equilibrium prices. Nevertheless, the results can be compared at the end of the experiment and provide information about the global outcome.

7. References

[1] Buyya, R., Economic-based Distributed Resource Management and Scheduling for Grid Computing. Ph.D. Thesis. Monash University, Melbourne, Australia, 2002. http://www.buyya.com/thesis/thesis.pdf.

[2] Caceres, R. et. al. , "Web Proxy Caching: The Devil is in the Details". *Proceedings of the Workshop on Internet Server Performance.*Madison, WI, 1998. http://www.research.att.com/~anja/feldmann/papers/wisp98_proxim.ps.

[3] Shankland, S. "Buddy, can you spare some processing time". *CNET News.com*. http://news.com.com/2100-1001-245266.html?legacy=cnet,