Modeling and Simulation of Media-On-Demand Services – Evaluating a Digital Media Grid Architecture

Sven Graupner, Winfried Kalfa[1], Carsten Reimann
Internet Systems and Storage Laboratory
HP Laboratories Palo Alto
HPL-2002-192
July 9th , 2002*

E-mail: {sven_graupner, carsten_reimann} @hp.com, winfried.kalfa@informatik.tu-chemnitz.de

multimedia content delivery, media grid, media centers, grid computing

Delivering digital media content over the Internet remains a challenging task. Work on Video-on-demand has focused on protocols and bandwidth allocation between content providers to consumers in the past. The Multimedia-Backbone (Mbone) has been a practical example. Systems like this focused on dealing with significant cross-network traffic, but not avoiding it. Content delivery and placement is now approaching the problem differently making use of the locality of groups of consumers. Popular media content can widely be shared and cached "at the edges" of the Internet close to where consumers generate demands. We propose a three-tier hierarchy of few initial media sources (original content providers) with distribution layers through regional and residential media centers acting as caches for popular media content. The behaviour of such an architecture is evaluated by simulation.

* Internal Accession Date Only Approved for External Publication
[1] Chemnitz University of Technology, Department of Computer Science, 09107 Chemnitz, Germany

Modeling and Simulation of Media-On-Demand Services – Evaluating a Digital Media Grid Architecture

Sven Graupner[1], Winfried Kalfa[2], Carsten Reimann[1]

[1)] Hewlett-Packard Laboratories, 1501 Page Mill Rd, Palo Alto, California, CA 94304, USA
e-mail: {sven_graupner,carsten_reimann}@hp.com
[2)] Chemnitz University of Technology, Department of Computer Science, 09107 Chemnitz, Germany
e-mail: winfried.kalfa@informatik.tu-chemnitz.de

Abstract: Delivering digital media content over the Internet remains a challenging task. Work on Video-on-demand has focused on protocols and bandwidth allocation between content providers to consumers in the past. The Multimedia-Backbone (Mbone) has been a practical example. Systems like this focused on dealing with significant cross-network traffic, but not avoiding it. Content delivery and placement is now approaching the problem differently making use of the locality of groups of consumers. Popular media content can widely be shared and cached "at the edges" of the Internet close to where consumers generate demands. We propose a three-tier hierarchy of few initial media sources (original content providers) with distribution layers through regional and residential media centers acting as caches for popular media content. The behaviour of such an architecture is evaluated by simulation.

Key words: media grid, media centers, multimedia content delivery, grid computing.

1. INTRODUCTION

Digital media content appears to be the major driving force for the further development of the Internet. Images and pictures have always been present in the World Wide Web, served as static files. Audio and music have become dominant when consumers started sharing music. Ignited by Napster and later continued in an distributed (or peered) fashion by Gnutella, digital music content dominates network traffic. AT&T, the largest American

Internet provider, states that more than 60% of its total traffic results from consumers swapping digital music content [1].

It appears to be natural that the next step will be digital video content distributed in various forms through the Internet. Movies are expected to be delivered through the Internet from the large distributors to end consumers. Similarly like music, consumers are expected to swap video content. Looking at volumes, there is roughly three orders of magnitude difference between pictures (range KB), audio/music (range MB), and video/movie (range GB). But video content is also different compared to pictures and audio in regard to the way it is being delivered. Pictures and audio are primarily delivered in a download-and-play/view fashion. Files are completely downloaded and can then be watched or played locally. Bandwidth and storage constraints mostly prevent this mode for digital video content today. Instead, only a short interval is usually pre-buffered along the path from the content provider to the consumer. Differently to off-line file downloads, video content delivery depends on the quality of the video stream. Quality of video content delivery is determined by the capacity of the server and the smallest-bandwidth along the path from the provider to the consumer [2]. The capacity of the server is usually specified in terms of the number of simultaneous transmissions (or channels). The bandwidth of the transmission path from the server to the consumer depends on a variety of factors beginning with the weakest bandwidth component, for instance a router or a network link, the overall delivery delay depending on the distance between server and consumer, the load conditions along the path, that is usually shared with other transmissions, and the variance or the jitter.

It is obvious that placing servers at the edges of the Internet close to where consumers generate demand for content has several advantages:

- Load reduction at the original content delivery servers by serving content from edge servers;
- Transmission paths between server and consumer can be kept short improving transmission quality;
- Overall network traffic reduction by serving demand locally and by thus improving the behaviour of the overall network;
- Cross-network traffic reduction.

In particular the last two aspects are very important for Internet Service Providers (ISP) and network operators. They have a strong interest that traffic remains local within their networks since they are being charged for traffic crossing into other providers' networks.

All these aspects foster the introduction of content caches at the edges of the Internet. However, traditional web caches cannot be used for video content delivery due to the specifics of video content (cannot be delivered as downloadable files). Streaming media servers are required for digital content delivery in form of multi-media streams. We introduce locations in consumer neighbourhoods where media servers are hosted as *media centers*.

In regard to the emerging large-scale computing, or Grid computing [3], the proposed architecture can provide a pattern for establishing a Grid for delivering digital media content to consumers, a grid named *Media Grid*.

We propose an architecture of a three-tier hierarchy of few initial media sources (original content providers), regional media centers and residential media centers acting as "caches" for popular media content. The architecture and its behaviour are evaluated by simulation. The simulation has been performed as a distributed simulation in a 524-node Linux cluster.

2. ARCHITECTURE

Content is initially induced into the system through one of the original content providers at the top of the hierarchy. Content then can be disseminated through the hierarchy by either being "pushed" down to regional or residential media centers on request (an example is making a brand-new movie available in underlying media centers before launch day), or content is "pulled" from residential or regional media centers when it is demanded from consumers, but not currently present.

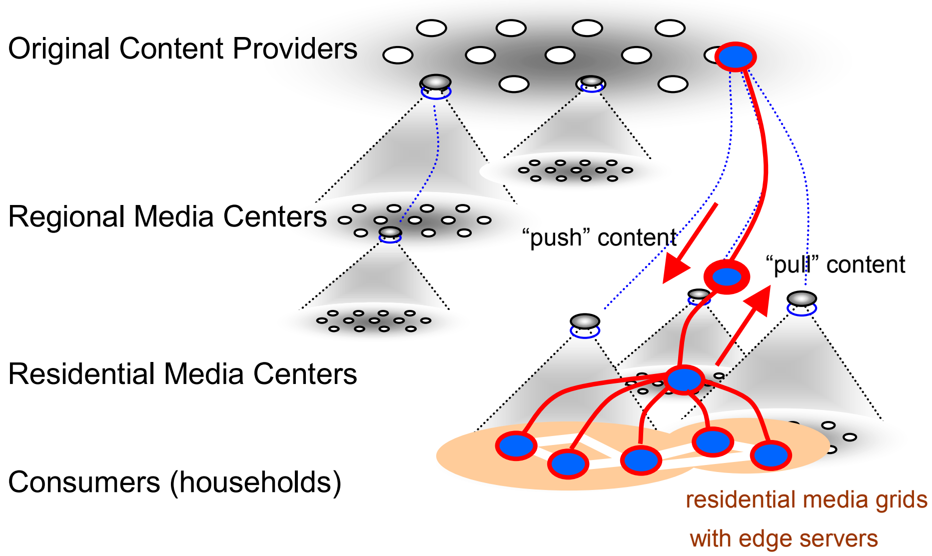

Figure 1. Architecture of the digital media content delivery Grid.

Names of the tiers have been chosen intentionally. Residential media centers are meant to be installed in consumers' neighbourhoods (at the inner end of the so-called "last mile"). Regional media centers are meant to aggregate a number of neighboured residential media centers, for instance of a metropolitan area. It is assumed that regional as well as residential media centers share the same temporal load pattern (located within one time zone).

Residential media centers are interconnected among each other according to their geographic neighbourhoods. Consumers are connected to one associated residential media center from which they obtain media content. On request of a consumer, and provided that the media content is available in the media center, the media content is directly served from the residential media center to the consumer's home. If content is not present, the residential media center asks its neighbour or peer residential media centers for the content. If the content cannot be obtained from there, the residential media center asks its associated regional media center. This procedure of asking peers before the higher-ordered element intends to avoid load at regional media centers.

Regional media centers behave similarly. If content that is asked for by assigned residential media centers is not available, regional media centers contact their peers and obtain content from there. If content cannot be obtained from peer regional media centers, content is obtained from the original content provider at the top of the distribution hierarchy.

Obtaining content always means the complete download of content and caching inside regional or residential media centers, not at the consumer site. Each media center has a certain capacity for storing content and maintains an index about the stored content. It is assumed that content can be uniquely identified by index data (keys), as well as that the original content provider can be derived from a content inquiry description issued by consumers.

3. SERVICE MODEL

The service model describes elements, relationships and parameters of the modelled environment. Elements are: consumers requesting media content, media centers, both residential and regional, and original content providers. For simplification, the two levels of media centers, regional and residential, have been consolidated into one level of media centers for the service model used for simulation. Relationships among elements are hierarchical as described before. Parameters are formulated based on queuing theory.

Consumers generate demand for the system. Consumer behaviour is modelled in three stages: inactive (no content requested), main time (occasional requests), and prime time (frequent requests). Content is classified into music (audio) and movie (video) content. Both are distinguished in terms of storage and transmission time.

The overall parameters are summarized:

M :	media content divided into music and movie content
$\frac{1}{\overline{\mu}^{1}_{n}}$	request processing time of media center *n* (time between an incoming request and the start of a transmission)
m_n :	storage capacity of media center *n*
$\_\ c_n$:	number of customers connected to media center *n*
$\eta_{n,m}$:	network capacity in KB/s from media center *n* to content provider *m*
$\mathrm{M}_m \subseteq \mathrm{M}$:	media content on content provider *m*
χ_m :	available transmission channels of content provider *m*

Parameters describing *media centers* are:

$\lambda^1_{0,n}(t,\Delta t)$:	rate of customer requests to media center *n*
$\mu^1_n(t,\Delta t)$:	request processing capacity of media center *n*
$\omega^1_{0,m}(t,\Delta t)$:	mean transmission time of requests to media center *n*
$\eta_{n,m}(t,\Delta t)$:	network traffic from media center *n* to content provider *m*
$\frac{1}{\mu^1_n(t,\Delta t)}$	mean time to process a request at media center *n*

Parameters describing *original content providers* are:

$\lambda^2_{n,m}(t,\Delta t)$:	request rate from media center *n* to content provider *m*
$\mu^2_m(t,\Delta t)$:	request processing capacity of content provider *m*
$\omega^2_{n,m}(t,\Delta t)$:	mean transmission time of requests to media center *n* to content provider *m*
$\frac{1}{\mu^2_m(t,\Delta t)}$	mean time to process a request at content provider *m*

Time progresses in the simulation in discrete steps of 2 minutes of real time. Each of the time steps triggers a series of events in the system, initiated by consumers requesting content. Initial events may trigger further events such as fetching content into a media center. For more detail on implementation as well as aspects of the distributed realization in a 524-node Linux cluster, see [4].

4. SIMULATION

In this section we describe exemplary the investigation of two scenarios to demonstrate how service models can be applied in the simulation system. For simplification, residential and regional media centers have been consolidated and referred to as media centers in the following. The simulation generates measurements behind the charts shown at the end of this section. In order to analyze the impact of caching media content in media centers, two scenarios have been considered:

- scenario 1: media access with or without the effect of media center caches,
- scenario 2: effect of preloading media content into media centers.

In both scenarios we use 3 content providers and 10 media centers with 2000 available media titles in total (1000 music and 1000 movie content). Content is randomly distributed across the 3 content providers. Each content provider stores 600 titles (300 music and 300 movie) exclusively, and 200 media titles (100 music and 100 movie) are shared. Content is accessible from all providers. We apply a data rate of 1.5 Mbps to movie content. The duration of a movie is assumed as 1.5 hours. For the 100 shared movies available from all content providers, we assume a duration between 15 and 30 minutes. The duration of music content is between 3 to 5 minutes. A data rate of 128 Kbps applies to music.

Each media center is connected to each content provider through a 100Mbps link. We assume 100 consumers connected to each media center. There are two groups of consumers: The first group accesses content between 08:00 p.m. and 10:30 p.m. The second group accesses content between 05:00 p.m. and 07:30 p.m. From the total of the 10 assumed media centers, each consumer group refers to 5 of them. Consumer groups were introduced to model consumers in two time zones generating demand with a time-shift.

We introduce popular (“top”) media titles for each consumer: One movie from specific content providers and three from the shared movies that are available from each content provider. Each content provider has 100 transmission channels and an average processing time of 0.5 seconds between the receipt of a request and the start of the transmission. A media center has a start-up time of 0.25 seconds per request.

4.1. Scenario 1: Media Access With or Without the Effect of Media Center Caches

4.1.1. Service Model Configuration

We can formulate the service model for the first scenario:

$\mathrm{M} = \{\{1,..,1000\}_{movie}, \{1,..,1000\}_{music}\}$: 1000 movie and 1000 music titles. Numbers identify movie and music titles in a range between 901 and 1000 representing the shared content available from providers.

$\frac{1}{\overline{\mu}_n^{1}} = 0{,}25\,\mathrm{sec}$ $\underline{c}_n = 100$ consumers connected to a media center
$\eta_{n,m} = 100 Mbps$ total transmission capacity

Content set M is partitioned into three sets $\mathrm{M}_{1,2,3}$, each representing a different category of content distribution:

- $\mathrm{M}_1 = \{\{1,..,300,901,..,1000\}_{movie}, \{1,..,300,901,..,1000\}_{music}\}$: 300 movie titles (1 to 300) that are exclusive for each provider and 200 shared movie titles (901 to 1000). The same ratio applies to music.
- $\mathrm{M}_2 = \{\{301,..,600,901,..,1000\}_{movie}, \{301,..,600,901,..,1000\}_{music}\}$ 300 movie titles (301 to 600) that are exclusive for each provider and 200 movie titles (901 to 1000) shared movie titles. The same ratio applies to music.
- $\mathrm{M}_3 = \{\{601,..,1000\}_{movie}, \{601,..,1000\}_{music}\}$: 300 movie titles (601 to 300) exclusively stored and 200 movie titles (901 to 1000) that are exclusive for each provider. The same ratio applies to music.

The individual processing behaviour of media centers for content sets is:

$\frac{1}{\overline{\mu}_m^{2}} = 0{,}5\,\mathrm{sec}$ $\chi_m = 100$ transmission channels of a content provider

The two scenarios differ in regard to storage capacity of media centers (measured in terms of titles that can be stored or cached):

$m_{n,movie} = 0; m_{n,music} = 0$: media center *n* in scenario 1,

$m_{n,movie} = 120; m_{n,music} = 120$: media center *n* in scenario 2.

For the simulation, we use the time interval between 00:00 a.m. and 11:59 p.m. with both categories of consumer behaviours.

4.1.2. Results

Both simulations for either scenario have been run three times and have been compared against each other. Since simulation results were nearly be the same, we refer to the results of the first simulation run for analysis.

Diagrams in *Figure 2* show the request trace appearing at content provider 1 ($\sum \lambda_{n,1}^{2}(t, 60\,\text{sec}); n \in \mathrm{N}, 1 \leq n \leq 10$). The dashed line shows the rate of refused requests:

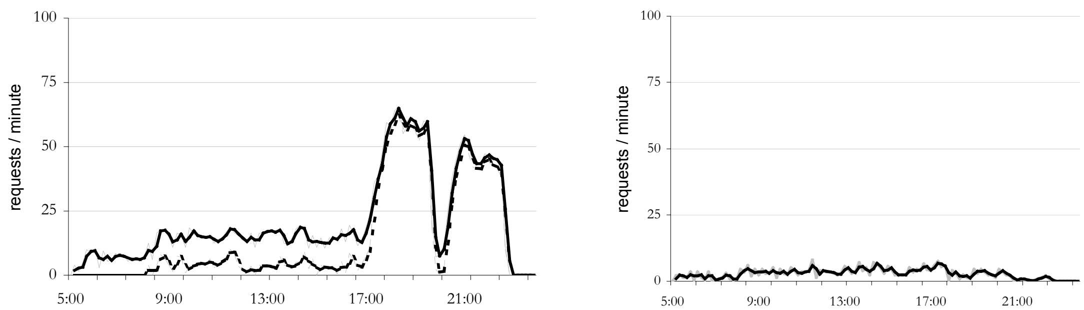

Figure 2. Request rates at content providers without (left) and with (right) media centers.

The left diagram shows two overload conditions between 05:00 p.m. and 00:00 a.m. The request rate appearing at content providers' sites shortly increased above 60 requests per minute causing the content provider to refuse incoming requests since it only has 100 transmission channels available with each channel being used for a duration of at least 3 minutes. The reason for the high request rate is that clients have been modelled to immediately repeat refused requests amplifying load conditions at content providers. Having intermediary caches available in form of media centers avoids this condition. This effect is shown in the chart at the right side.

Diagrams in *Figure 3* show the network traffic exemplary observed between media center 1 and content provider 1 ($\eta_{1,1}(t, 1\,\text{sec})$):

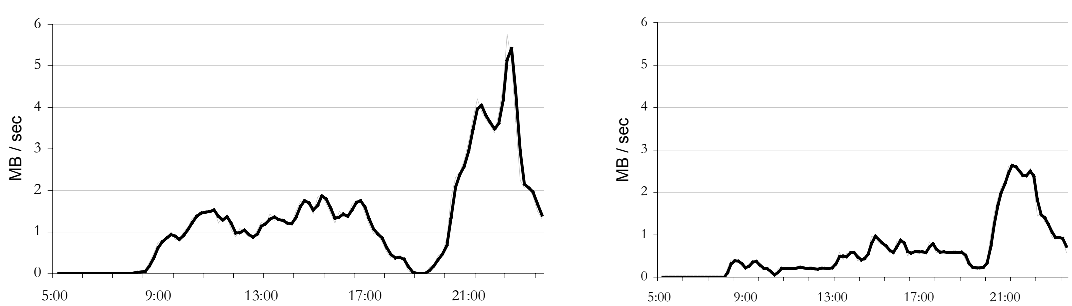

Figure 3. Network load without (left) and with (right) the effect of media centers.

Diagrams in *Figure 3* show, that using the cache effect of media centers reduces traffic in the network. There are even periods with no traffic at all (between 01:00 p.m. and 02:00 p.m.) when all content could be served from local media centers caches. *Figure 4* summarizes that media centers significantly reduce network traffic, in the shown case (traffic observed at content provider sites) by nearly 48% compared to the case without media centers.

from	Without media centers (KB)	With media centers (KB)
Content provider 1	84,899,393	41,848,576
Content provider 2	86,603,076	35,311,162
Content provider 3	72,505,747	39,783,190
total	**244,008,216**	**116,942,928**

Figure 4: Total network traffic observed at content provider sites.

Conclusion: As expected, it could be shown that media centers avoid hot spots at content providers and significantly reduce network traffic.

4.2. Scenario 2: Effect of Preloading Media Content to Media Centers

In the second scenario, we examine the case when a new movie is being released and should be launched at the same time throughout the system. What is the effect of preloading content to media centers before launch date?

4.2.1. Service Model Configuration

In order to investigate this scenario, we assume a simple example of only one content provider and one media center. The configuration parameters are alike the parameters of the first example, now with a single content provider:

$\mathrm{M} = \{\{1,..,1000\}_{movie}, \{1,..,1000\}_{music}\}$: 1000 movies, 1000 music titles.

$\mathrm{M}_1 = \{\{1,..,1000\}_{movie}, \{1,..,1000\}_{music}\}$: All media titles are stored at content provider 1.

$m_{1,movie} = 120; m_{1,music} = 120$: storage capacity of media center 1.

For the simulation, we use the same time interval between 00:00 a.m. and 11:59 p.m. with a measuring period of $\Delta t = 2\,\mathrm{min}$. All new movies have been preloaded into all media centers.

4.2.2. Results

When consumers wanted to watch the same movie, the associated media center requests the media title from a content provider only once (after the first request from a consumer). The following request then is served from the local media center cache avoiding further requests to a content provider. Diagrams in *Figure 5* show that preloading media content has no visible impact on request rates observed at content providers, which is not intuitive.

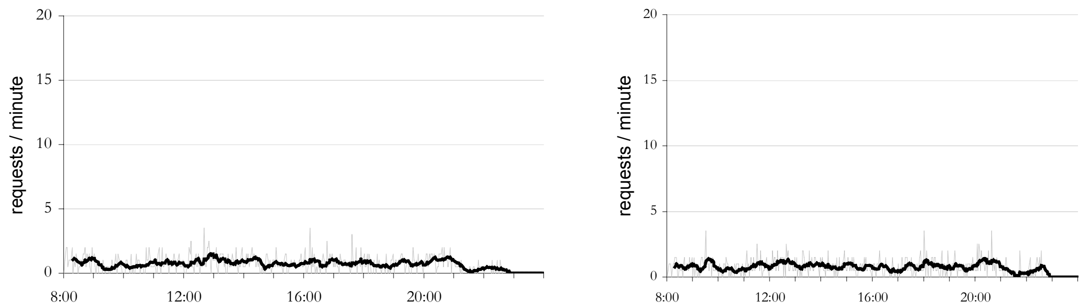

Figure 5. Request rates at content providers with (right) and without (right) preloading.

This result is similar in regard to the observed network traffic: the effect of preloading one movie is negligible.

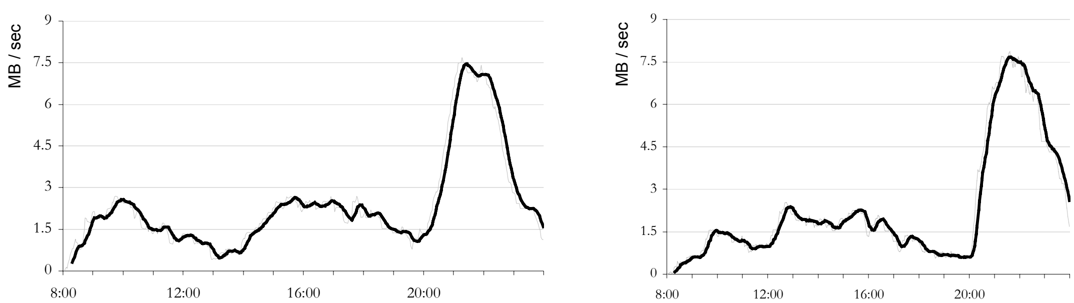

Figure 6. Network traffic without (left) and with (right) preloading media content.

Conclusion: Preloading media content seems to have no influence in the shown scenario. This can be explained by the effect that only one movie has been chosen for launch at a particular date with an unchanged consumer behaviour for watching other movies. The demand for this one new movie can thus be neglected compared to the other content requested by consumers. Further investigations (not included here) show that preloading content has an advantage when consumers simultaneously and in large numbers chose watching the new movie at launch date. Consumer behaviour is then different with large numbers of consumers asking for the same, new movie. In this situation, caches in media centers can absorb the peak demand that would have been occurring at content providers for a movie at launch date.

5. RELATED WORK

The architecture imposes a hierarchical order of its elements in contrast to other content delivery systems such as Content Addressable Networks (CAN) [5] that allow a rather loose organization of peered locations where content can be stored in a virtual address space. Content inquiries (keys) are mapped into coordinates of locations where content can be found. This organization is useful for the very bottom layer of consumers directly exchanging content (Napster, Gnutella and other p2p systems fall into this space). It is hard to maintain a hierarchical structure in such an environment with a high degree of fluctuation and dynamism. It is suitable for end-consumer environments. We rather address the "retail layer" of the distribution system, the layer above consumers that is usually transparent to them and under the control of the ISP or network operator. Here, hierarchy can be imposed and is easier to maintain than loose organizations of p2p systems. Akamai [6] is an example for delivering web content through a hierarchical cache organisation to consumers very much alike we have described here for delivering digital media content. The convergence of web and digital media content is foreseeable with web content distributors moving towards also distributing digital media content.

In [7], an architecture for distributing multi-media content in an entire peer-to-peer fashion is proposed based on a network of Multiple Independent Indexing Servers. This approach also addresses the "retail" layer above consumers, however it assumes a peer coupling of nodes, not a hierarchy.

6. SUMMARY AND FUTURE WORK

We have proposed an architecture of content providers, regional and residential media centers for distributing media content. Simulation results have shown that network traffic can be avoided when requested content is served locally from media centers rather than directly from original content provider sites. Overload conditions at content providers can be avoided, and the overall network traffic can be reduced.

Media centers are expected to make use of specific capabilities Utility Data Centers provide [8], [9]. They will allow to establish an autonomous service control system for delivering media content to consumers with the specific capability to migrate cached content and services automatically between media centers following demands [10]. This describes our future research direction.

REFERENCES

[1] Chawathe, Y.: Can Heterogeneity Make Gnutella Scalable?, Stanford Distributed Systems Seminar, http://www.stanford.edu/class/cs548/schedule.shtml, May 1, 2002.

[2] Carter, S.R., Paris, J.F., Mohan, S., Long, D.D.E.: A Dynamic Heuristic Broadcasting Protocol for Video-On-Demand, in Proceedings of the 21st IEEE International Conference on Distributed Computing Systems, Phoenix, AZ, May 2001.

[3] Foster, I., Kesselman, C., Nick, J.M., Tuecke, S., The Physiology of the Grid – An Open Grid Services Architecture for Distributed Systems Integration, DRAFT, http://www.globus.org/research/papers/ogsa.pdf, May 2002.

[4] Reimann, C.: Model-Based Monitoring in Large-Scale Distributed Systems, Master's Thesis, Chemnitz University of Technology, Department of Computer Science, May 2002.

[5] Ratnasamy, S., Francis, P., Handley, M., Karp, R., Shenker, S., A Scalable Content-Addressable Network, SIGCOMM 2001, San Diego, August 27-31, 2001.

[6] Meggs, B.: Global Internet Content Delivery, talk at the Computer Science Colloquium, Stanford University, April 9, 2002, Akamai Technologies Inc., http://www.akamai.com.

[7] Kalogeraki, V., Delis, A., Gunopulos, D.: Handling Multimedia Objects in Peer-to-Peer Networks, pp. 438-439, Proceedings of the 2nd IEEE/ACM Symposium on Cluster Computing and the Grid CCGrid 2002, Berlin, Germany, 21-24 May, 2002.

[8] *HP Utility Data Center*, http://www.hp.com/go/hpudc, http://www.hp.com/go/always-on, November 2001.

[9] Graupner, S., Kotov, V., Trinks, H.: Resource-Sharing and Service Deployment in Virtual Data Centers, IEEE Workshop on Resource Sharing in Massively Distributed Systems (ICDCS-2002), July 2, 2002, Vienna, Austria.

[10] Andrzejak, A., Graupner, S., Kotov, V., Trinks, H.: Self-Organising Control in Planetary-Scale Computing, pp. 359-364, Proceedings of the 2nd IEEE/ACM Symposium on Cluster Computing and the Grid CCGrid 2002, Berlin, Germany, 21-24 May, 2002.

[11] Graupner, S., Kotov, V., Trinks, H.: *A Framework for Analyzing and Organizing Complex Systems*, Proceedings of the 7th IEEE International Conference on Engineering of Complex Computer Systems (ICECCS 2001), pp. 155-165, Skövde, Sweden, June 11-13, 2001.