Content Distribution Infrastructures

Carsten Griwodz, Thomas Plagemann, Ralf Steinmetz

2nd July 2004

1 Public Outreach

Since the early days of the world wide web (WWW), the information infrastructure provided over the Internet has improved considerably. The simplicity of offering and accessing data on the WWW and the increase of commercial uses of the Internet are major reasons for this development. During this development, it was discovered that shifts in popularity of content and services offered over the web lead to unbalanced load. A sudden increase in popularity can occur (slashdot effect) and is very hard to predict. Therefore it is also very hard to scale these servers and their connection to the Internet at the right time. To overcome these problems, hierarchies of cache servers such as harvest and squid have been set up to alleviated general load, but the large amount of content has limited their effectiveness in many situations. For that reason, controlled pre-distribution of content on behalf of the content providers has been offered as commercial service. To offer the service, servers are installed in the networks of different Internet service providers around the world. The servers can cooperate but they are only connected through the Internet. An infrastructure for this kind of service is known as content distribution network (CDN). Offering content through such a CDN has two advantages. First, several providers' content is offered from the same CDN and only the content of very few providers will experience a sudden increase in popularity at the same time. By offering the content through a CDN, not every content provider must maintain servers that can cope with the load that he experiences at times of very high popularity. Instead, the CDN hosting the content will experienced a multiplexed load, and it can be scaled to cope with high load for some but not all content providers. Second, content in the CDN may be replicated over several hosting centers and accesses to content can be redirected to the closest one, thus reducing network load and access latency.

By now, these commercial content hosting CDNs that offer access to discrete media constitute the classical form kind of CDNs. There are, however, other forms of content distribution that are covered by the term CDN as well. One form is used to distribute a different kind of content, namely live data streams that are transmitted from a single source to a large audience. Another form distributes discrete media as well, but uses the peer-to-peer (P2P) model in which equal nodes collaborate for the distribution, and in which nodes are often owned by individuals. Even though commercial CDNs for live and stored content and P2P systems have been investigated separately so far, they share characteristics, for example that they consist of nodes that are connected through the Internet but typically not directly to each other. The integration of these approaches is therefore a challenge to researchers. Another is the increased importance of multimedia content and interactive applications, which will impose new problems on CDNs.

This seminar has brought together researchers who address the challenges that lie in the improvement of CDNs. Among the challenges faced are those related to scaling of the CDN infrastructure, the use of appropriate techniques and tools, and the management of the growing CDNs. Several attendees of the seminar presented ideas and results concernings this topic. Topics covered in presentations of participants' current research include system and network support for scalability of content and CDN nodes, the communication between nodes and from the nodes to end-systems. The topics included dimensioning, scaling, configuration and reconfiguration of distribution networks for both, hierarchical CDNs and those that follow the P2P model, and reports on the needs of variations applications that rely on CDNs now or in the future.

Dagstuhl Seminar Proceedings 04201
Content Distribution Infrastructures
http://drops.dagstuhl.de/opus/volltexte/2006/502

It became clear that the interaction of end users with services offered via the CDN will become highly important in the future, and that current CDNs can not address this demand in an appropriate manner. Interactivity was therefore identified as a crucial point for the future development of CDNs. Related to this point are danger of denial-of-service attacks on CDNs and questions about the most appropriate support for heterogenous devices and the distribution of applications rather than content.

2 Scientific Highlights

The CDN seminar was meant to inspire active discussion among participants, and to achieve this, a major share of the time was set out for stating opinions and for the discussion of future problems in working groups. The seminar was therefore structured into time for presentations of recent and ongoing work of researchers and PhD students, a session for the presentation of so-called outrageous opinions, and the group work.

2.1 Presentations

In their presentations, the participants presented their recent results and ongoing work related to CDNs. A topic addressed by several was the scalability of CDNs. Both **Nikos Laoutaris** and **Adrian Cahill** work on the dimensioning of hierarchical CDNs, and the approaches that they presented are aimed at the replication of complete objects to the nodes of a CDN. Nikos approaches this by finding the most appropriate location of complete content objects and determining node placement and object placement concurrently. He argues that such dimensioning can be beneficial for offline and online dimensioning of CDNs. The first is useful for planning of CDNs, the second in cases where resources are not fixed and can be rented, for example. Adrian's goal is similar, but he considers an applications that focusses on the distribution streaming media rather than discrete content. He is also concerned with appropriate online clustering of requests to find appropriate nodes for groups of clients with similar interests. The assignment of partial rather than complete objects, and the exploitation of presentation structure was the topic of **Frank Johnsen**'s presentation. He demonstrated that a new kind of partial caching for stored multimedia content, called structured partial caching, has the potential to improve the performance gains of full objects caching, as well as temporal partial and quality partial caching schemes. The aforementioned three participants state that less popular content should be stored in fewer places in the CDN, and this it would have to be delivered over longer distances. For this kind of content, but also when live content is distributed or when scalable delivery protocol place relevent further away from end-systems, appropriate delivery over longer distances is required. For systems where this delivery is handled within the CDN, **Trude Hafsøe**'s work comes in. She presented arguments for the reconfiguration of overlay networks that are formed by the nodes of a CDN. The idea is that CDNs should be aware of the quality of the paths between their nodes, in order to choose the most appropriate sequence of nodes (an overlay route) for the delivery of content through the CDN. Based on existing measurements, she identified three types of changes in network quality, two of which can be addressed by reconfiguration of overlay routes.

While the presentations mentioned above were mostly concerned with CDNs that are hierarchically organized or at least organized in a controlled manner, several participants were concerned with the use of P2P approaches for content distribution. **Laurent Mathy** challenged the attendees with the demand for autonomic CDNs that combine classical CDN approaches with ideas taken from P2P systems. Such autonomic CDNs should manage the distribution of content in the CDN autonomously to maintain a good performance and a high degree of availability, and to survive without human intervention in case of failures. **Thomas Plagemann**, **Pascal Felber** and **Jochen Mundinger** showed how they address these challenges. Thomas Plagemann proposed the use of a hybrid CDN that makes use of ideas from the BitTorrent P2P system. He presented an analysis of BitTorrent that demonstrated how firewalls inhibit the effectiveness of contributing to resource sharing in the BitTorrent system. To avoid these impediments, he proposed the use of cache servers that can contribute on behalf of their clients and improve the scalability of the cache and main server. Pascal Felber showed how the effectivity of content distribution to a large number of receiver nodes can be improved considerably by using P2P concepts instead of hierarchical distribution. In every system that distributes content that is large enough to consume a relevant share of network or system resources on the original server, the contribution of subordinate nodes (or peers in a P2P system) can reduce the load per node considerable and speed up the distribution process. He shows the advantage of a meshed distribution approach over a hierarchical one. A very similar goal is followed by Jochen Mundinger, but he focusses on the problem that participants in the distribution process have limited uplink capacities. In contrast to the organization in tree structures proposed by Pascal Felber, his approach relies on availability lists that are provided by directory servers and that allow nodes to select randomly from the set of nodes that can be contacted to download missing parts of a file.

In one part of this talk, **Dwight Makaroff** addressed another option for improving the scalability of distribution systems, scalable delivery protocols. Using these protocols, content can be partitioned into segments that may be transmitted out of order and at different speeds than playout speed, using multicast and broadcast techniques. Combinations of using these protocols with both hierarchical CDNs and P2P systems have been presented in the past. And he pointed out that an improvement of the distribution is not sufficient but that system performance must also be taken into consideration. Concerning this issue, he addressed issues of server resources, in particular disk admission control. As a complementary means of ensuring the scalability of distribution systems, **Charles Krasic** presented and demonstrated the use of quality-adaptive content in distribution systems, which makes it possible to adapt to available network and end-system resources. This approach is particularly important when the delivery of content from CDN nodes to end-systems is affected by other Internet traffic, or when heterogeneous devices are served from the same content object. **Jens Brandt**[1] follows an approach to system scalability that is an alternative to Charles Krasic's. He goes a step towards CDNs that support applications, and proposes transcoding modules in the nodes of CDNs. Using these modules, he achieves the adaptation of streams to the capabilities of end-systems, in particular mobile nodes that may lack the processing power for presenting data streams that are originally stored in the CDN.

While the participants mentioned so far emphasize classical applications for content distribution, that are aimed at the delivery of stored content or the distribution of live media streams, **Verena Kahmann**, **Jörg Widmer** and **Michael Zink** proposed new applications that have not attracted so much attention yet. Verena Kahmann claims that media streaming applications in which multiple users and their devices cooperate will appear, for example in home networks, learning environments and spontaneous meeting scenarios. She showed how IETF protocols can be used to build an infrastructure for such collaborative media streaming applications. This approach constitutes a means for merging several kinds of applications that have been addressed separately and in proprietary manners in the past. The application that Michael Zink is concerned with is more unusual with respect to existing work on CDN, because it requires the collection of data from high-volume data streams rather than the distribution of content. The `casa` project collects weather radar data at several stations and must transfer this data over shared data links to a central archival and computing center. This project demands the adaptation to variable network capacity, overlay routing, and scheduling of content for transmission. An application that imposes yet other demands on CDNs supporting it are networked multiplayer games, and Jörg Widmer introduced how this kind of applications can draw benefits from application support in CDNs. Games are among the applications that require live distribution, but in contrast to established uses of CDNs, this application area is stamped by frequent changes to priority with which nodes must communicate with each other. Therefore, he proposes a variation of a spanning tree approach for overlay routing that takes application semantics into account.

2.2 Outrageous opinions

Some participants volunteered to present a short, controversial statements to stimulate discussions during the workshop and in the working groups in particular. **Martin Karsten** reminded us of the need for quality of service support, and that the need for maintaining state in the network is especially important

[1] attended under the name Jens Zechlin

for large-scale distribution. He pointed also out that, in contrast to public lore, statefulness in disguise is already in place in the Internet. His examples were mobile IP and IP multicast. **Oliver Spatscheck** asked critical questions about the usefulness of the P2P paradigm. He provided arguments for the assumption that centralized hosting would improve availability and resource efficiency over the typical, decentralized P2P model.

Susanne Boll, **Klara Nahrstedt** and **Wolfgang Leister** presented outrageous opinions that were more immediately concerned with future applications that may have relevance for the development of CDNs. They shared the opinion that the support of personalization will become a major challenge for the future. Susanne Boll observed a trend towards mixed media and mixing of life and stored content that must be delivered over heterogenous networks to different kinds of devices. Users will demand more personalized content and therefore broadcasting multimedia content in combination with a suitable return channel will be a more important application scenario than streaming live content to a large audience that are connected via point-to-point connections. Similarly, Wolfgang Leister stated his belief that the importance of personalized content will increase. He proposed that more content will be rendered on end-systems from data streams rather than being stored in a form that is easily played out. In contrast to Susanne Boll, he claimed that this personalization will lead to an increased relevance of communication between small groups of users, which is favoured by a P2P model and less by broadcast models. Klara Nahrstedt extended these predictions and claimed four major points: that more content will be generated by many difference sources, that media content will be produced in a distributed manner, that content aggregation through approaches such as mixing and summarization will play an important role, and that interactivity and customization to single users and groups will be required. Therefore, she defined the hourglass infrastructure for content distribution. This is based on the assumption that personalization and customization should be approached by collecting input from several sources and delivering output to several destinations, but that the data would pass through a single content aggregation point for each user or group that is meant to receive a particular customization.

The outrageous opinions were presented before the selection of topics for working groups, and it they did have major influence on the discussions in those groups.

2.3 Group work

The participants were presented with a couple of potential work group topics, from which the following four were chosen as topics:

- Future Architectures,
- P2P and CDNs,
- Future Bottlenecks, and
- Future Applications.

The future architectures working group predicted that we will see a variety of different types of CDNs in the future. These will be similar to current file sharing P2P systems or provide services similar to Akamai, but within these types, the provided services will be improved. Future CDNs will be very structured and well-managed, so that they can offer a fairly constant quality-of-service to their customers. These services will be offered as service level agreements. This will happen even though we can not expect end-to-end quality-of-service guarantees that are offered to end users. However, these service level agreements will include services such as transcoding of content in the CDN and guarantees concerning security. The working group considered that these CDNs would distribute code and software as well as content. They observed that we have the pieces for such CDNs available already today, but that they must be integrated in order to build multiservice CDNs, and to enable setting up the network, the content, transmission paths in the network, and management of the CDN. The group predicted that tools for customization are coming, but they indicated a lack of knowledge about application domains.

The future applications working group, on the other hand, identified future important application areas. These application areas will not all be covered by one kind of CDN, but different types of CDNs will exist for different applications. They will be distinguished by the applications' qualitative and quantitative requirements. One of the most important distinctions between these types will be their financial goal: whether they are deployed to earn money or to save money. The working group predicted that future applications will require the distribution of structured information and true multimedia documents, where media are combined and depend on each other. Future applications will use stored and live content together. Streamed data will play an important role, and support for interactivity will be required. To make this possible, CDNs can not only provide support for data but they must also support handling of metadata. The CDN must provide the means for automatic generation and support of meta data, its storage and distribution. CDNs can ignore neither security considerations not intellectual property rights issues.

The P2P in CDNs working group identified what P2P systems can contribute to CDNs. They identified benefits of P2P systems, such as decentralization, autonomy of nodes, the ability for self-organization of their distribution infrastructure, and the cooperation among nodes. It was also found that P2P systems share more than just content, as classical CDNs do. In addition to the distribution of content, P2P systems allow for sharing of CPU cycle, chat-like communication, sharing of storage and bandwidth. To determine which advantages and disadvantages the use of P2P approach would have when use in CDNs, the working group distinguished between node between communication between nodes of the CDN, and between the CDN and end systems. For communication among CDN nodes, they determined that the self-organization of P2P systems, their adaptivity, autonomic behaviour, their ability to react to strongly bursty requests would be beneficial, and that P2P ideas could probably be used to reduce maintenance cost for CDNs. However, protocols and interactions in P2P systems are more complicated, and control over the system gets lost. For communication between CDN nodes and end systems, they see the advantages that load is offloaded from the CDN and therefore, costs can be reduced. Also in this case, bursty access can be provided for. However, quality-of-service, reliability and availability guarantees can not be given, and the load on end systems is increased. They do therefore propose to explore hybrid solutions that can for example exploit proximity. The groups warned that because of the way in which P2P systems distributes content, it does not provide an appropriate solution for unpopular content.

The bottlenecks working group determined that bottlenecks come into play only if a CDN serves a massive number of users, offers massive amounts of dynamically generated data, or both. The group came to the conclusion that bottlenecks, that are usually considered bad, are not necessarily problematic if the CDN that experiences the bottleneck is a commercial one. First of all, many bottlenecks can be solved by investing in upgrades to nodes and networks of the CDN. Such upgrades are only beneficial when they are affordable as well. In existing CDNs, customers are already now asking for cheaper service rather than better perfomance, which indicates that the remaining bottlenecks for the delivery of static content are acceptable to paying customers of the CDN. However, upgrades depend on problems that can be parallalized efficiently, and therefore, services that require serial processing, or that can not be parallalized at acceptable cost form that most relevant bottlenecks of the future. This kind of bottleneck appears when content is dynamic, when it is personalized, and when users interact with the system. Even nowadays it is the case that CDNs are used extensively for the distribution of static content, while dynamic content is served from CDN customers directly, even if static and dynamic elements appear on the same web page. Future challenges concerning bottlenecks are therefore the reduction of those parts of applications that are transactional and can not be parallelized, and the resistance of CDNs against denial-of-service attacks. It is furthermore important to determine realistic workloads for CDN research, and work on the interoperation of different CDNs to avoid bottlenecks that shift dynamically due to autonomous reconfiguration.

For the last day, all working groups were inspired to contribute thoughts on the development of a research CDN. It became clear that two different demands exist that would best be solved by establishing two CDNs. The first demand is that for a CDN for teachers and researchers that can be used to exchange information. The second demand is for an infrastructure for realistic experiments with CDNs. While it might seems like a good idea to use experimental CDNs for the teaching and research material, the availability of such a CDN would not be acceptable.

Therefore we can distinguish the two goals. A CDN for researchers would contain educational content, research data, publications, and results. It should be cheap, easy to maintain with small effort, use standard technologies and be based on standards. It must support metadata, location of content, interactivity and live as well as stored content. It should also be integrated with groupware technology. For such a CDN, design choices would have to be made once, concerning centralized versus decentralized organization, planned versus on-demand replication, and the required fault tolerance and availability. Once the CDN is in place, it could be used for measurements and logging, and could provide realistic workload for experimental CDNs. A starting point for such a CDN could be BitTorrent, even though variations such as global resource sharing and the incentive to contribute to the distribution of undesired content might be required.

Such experimental CDNs should allow all participants to work within their own area of interest. This should be possible in parallel, with little influence on each other. It was therefore proposed to interconnect a limited number of physical nodes that implement a larger number of virtual CDN nodes each. These virtual nodes could be monitored in detail, such that measurements at the IP level, OS level, application level and even logging of hardware operations would become possible. A CDN toolkit should provide for services to simplify experiments, and provide monitoring tools as well. To allow for flexibility in the deployment of various experiments, the components of the CDN toolkit should be pluggable and allow the replacement of high as well as low level mechanisms. This would include functions for application-level addressing, indexing, searching, data placement, replication, management, cooperation and coordination. Since the amount of control required over this kind of research CDN, it was proposed that a centralized system with emulated networking such as Emulab may be a more appropriate approach than actual physical decentralization.

3 Perspectives

In the workshop, have discussed the future of CDNs in research and pratice. We were presented with current CDN research performed by the participants, and identified several challenges now and in the future. The current challenges are clearly, scalability and management of CDNs for stored and live content (in addition security, which was outside the seminar's focus). Future challenges arise from new applications, ownership models, and threats to the operation of CDNs.

We have identified a trend for demanding application support from CDNs. This means that we face the challenges of supporting interactivity, metadata handling, dynamic content generation, scaling and adaptation in future CDNs. CDN customers will demand QoS guarantees, while competition will require tradeoffs between price and performance. Content will in the future be more personalized than today, and we have also seen that future CDNs will in a much larger degree than today receive content from end users rather than just deliver it to them. To support all of these demands, future CDNs will much more than today support distributed application rather than only distributing content. The dynamicity and interactivity of future applications will become a major problem for all those applications that are not completely parallelizable, and pose a variety of research issues.

4 Keywords

CDN, P2P, Research CDN