

Content Distribution over IP: Developments and Challenges

Adrian Popescu, Markus Fiedler

*Blekinge Institute of Technology
SE-37179 Karlskrona, Sweden*

Demetres Kouvatsos

*University of Bradford
Bradford, West Yorkshire BD7 1DP, UK*

The telecommunication industry is actually facing two serious challenges with implications on future architectural solutions. The first challenge is regarding the irreversible move towards IP-based networking. The second challenge is regarding the deployment of broadband access in the form of diverse Digital Subscriber Line (DSL) technologies based on optical fiber and high-capacity cable but also the WiMAX access (IEEE 802.16 Worldwide Interoperability for Microwave Access) to provide high bandwidth access to home networks as well as to small and medium-sized businesses. Altogether, these developments offer the opportunity for more advanced and more bandwidth-demanding multimedia applications and services, e.g., Internet Protocol Television (IPTV), Voice over IP (VoIP), online gaming. A plethora of QoS requirements and facilities are associated with these applications, e.g., multicast facilities, high bandwidth, low delay/jitter, low packet loss. Furthermore, a very important issue is regarding the perceived QoS and the standards associated, e.g., as defined in ITU-T BT.500.11, ITU-T P.862. Even more difficult is for the service provider to develop a networking concept and to deploy an infrastructure able to provide end-to-end (e2e) QoS for applications with completely different QoS needs. On top of that, the architectural solution must be a unified one, which is independent of the access network and content management (i.e. content acquisition, storage and delivery). Other facilities like billing and authentication must be provided as well.

The foundation of multimedia distribution is provided by several components, the most important ones are services, content distribution chain, protocols and standards. The fundamental idea is to use the Internet for content acquisition, creation, management and delivery. Furthermore, an important goal is to offer the end user the so-called Triple Play, which means grouping together Internet access, TV and telephone service into one subscription on a broadband connection. Other important issues are billing and content protection, e.g., copyright issues, encryption and authentication (Digital Rights Management).

The convergence between fixed and mobile services that is actually happening in the wide and local area networking is expected to happen in the home networking as well. This puts an extra burden on multimedia distribution, which means that wireless access solutions of different types (e.g., WiMAX) must be considered as well. The consequence of throwing Triple Play into wireless services is the upcome of Quadruple Play.

It is therefore important to consider mechanisms and protocols put forth by the Internet Engineering Task Force (IETF) to provide a robust and systematic design of the basic infrastructure, and protocols like Session Initiation Protocol (SIP), IP DiffServ (RFC 2474/2475), together with Multi Protocol Label Switching (MPLS) and traffic engineering (RFC 3031), should be taken into consideration as possible solutions for the QoS control in core networks. Another important IETF initiative is regarding content distribution issues, which are addressed in the IETF WG for Content Distribution Networks (CDN) and Content Distribution Internetworking (CDI). Furthermore, new developments within wireless communications like IP Multimedia Subsystem (IMS) are highly relevant for such purposes. Similarly, the new paradigms recently developed for content delivery application-based routing (e.g., based on Peer-to-Peer (P2P) solutions) can be considered as alternative solutions for the provision of QoS on an e2e basis, without the need to replace the IPv4 routers with IP DiffServ routers. The main challenge therefore is to develop an open architectural solution that is technically feasible, open for future updating and services and cost-effective.

1 Content Distribution Networks

Content Distribution Networks (CDNs) are networking solutions where high-layer network intelligence is used to improve the performance in delivering media content over the Internet, e.g., static or transaction-based Web content, streaming media, real-time video, radio. The fundamental concept is based on distributing content to

cache servers located close to end users, resulting so in better performance, e.g., maximize bandwidth, minimize content latency and jitter, improve accessibility. CDNs are composed by multiple Points of Presence (PoP) with clusters (so-called surrogate servers) that maintain copies of (identical) content, resulting so in better balance between cost for content providers and QoS for customers. CDN nodes are deployed in multiple locations, in most cases placed in different backbones. They cooperate with each other, transparently moving content so as to optimize the delivery process and to provide users the most current content. The optimization process may result, e.g., in reducing the bandwidth cost, improving availability and improving QoS.

The client-server communication flow is replaced in CDN by two communication flows, namely between the origin server and the surrogate server and between the surrogate server and the client. On top of that, questions related to QoS, content multicasting and multipath routing heavily complicate the picture. Requests for content delivery are intelligently directed to nodes that are optimal with reference to some parameter of interest, e.g., minimum number of hops, or networks, away from the requester.

Organizations offering content to geographically distributed clients usually sign a contract with a CDN provider and distribute the content over the selected CDN by using a specific overlay model. Today, some of the most popular commercial CDN providers are Akamai, Nexus, Mirror Image Internet and LimeLight Network.

It is also important to mention that content distribution can be done by using IP multicast as well. In such a case, specific code is deployed in IP routers or switches such as they are able to recognize specific application types and make forward decision of their own that are based on predefined policies.

In practice, there are several challenges that must be solved in order to offer high-quality distribution at reasonable prices. Some of the most important questions are related to where to place the surrogate servers, which content to outsource, which practice to use for the selected content outsourcing, how to exploit data mining over CDN to improve the performance and what model to use for CDN pricing.

For instance, it is very important to choose the best network placement for surrogate servers since this is critical for the content outsourcing performance. A good placement solution may also have other positive effects, e.g., by reducing the number of surrogate servers needed to cover a specific CDN. Actually, several placement algorithms have been suggested, e.g., Greedy, Hot Spot and Tree-Based Replica, each of them with own advantages and drawbacks.

Another challenge is the selection of the content that should be outsourced in order to meet the customers needs. An adequate management strategy for content outsourcing should consider grouping the content based on correlation figures or access frequency and replicate objects in units of content clusters. Furthermore, given a specific CDN infrastructure with a given set of surrogate servers and selected content for delivery, it is important to select an adequate policy for content outsourcing, e.g., cooperative push-based, uncooperative pull-based, cooperative pull-based. These policies are associated with different advantages and drawbacks, today however most of the commercial CDN providers (Akamai, Mirror Image) use uncooperative pulling. This is done in spite of non-optimal solutions used to select the optimal server from which to serve the content. The challenge is to provide an optimal trade-off between cost and user satisfaction and new techniques like caching, content personalization and data mining can be used to improve the QoS and performance of CDN.

An important parameter to be considered is related to the CDN pricing. Today, some of the most significant factors affecting the pricing of CDN services are bandwidth cost, traffic variations, size of content replicated over surrogate servers, number of surrogate servers, and security cost associated with outsourcing content delivery. It is well known that cost reduction occurs when technology investments allow for delivering services with fewer resources. The situation is however more complex in the case of CDN since higher bandwidth and lower bandwidth cost also have as a side effect that customers develop more and more resource-demanding applications with harder and harder demands for QoS guarantees.

The fundamental entities of a CDN are the network infrastructure, content management, content routing and performance measurement. Content management is about the entire content workflow, from media encoding and indexing to content delivery at edges including also ways to secure and manage the content. On the other hand, content routing is about delivering the content from the most appropriate server to the client requesting for it. Finally, performance measurement is considered as part of network management and it is regarding measurement technologies used to measure the performance of the CDN as a whole.

2 Routing in Overlay Networks

Overlay networks have recently emerged as a viable solution to the problem of content distribution with multicasting and QoS facilities. Overlay networks are networks operating on the inter-domain level, where the edge hosts learn of each other and, based on knowledge of underlying network performance, they form loosely

coupled neighboring relationships. These relationships are used to induce a specific graph, where nodes are representing hosts and edges are representing neighboring relationships. Graph abstraction and the associated graph theory can be further used to formulate routing algorithms on overlay networks. The main advantage of overlay networks is that they offer the possibility to augment the IP routing as well as the Quality of Service (QoS) functionality of the Internet.

One can state that, generally, every P2P network has an overlay network at the core, which is mostly based on TCP or HTTP connections. The consequence is that the overlay and the physical network can be completely separated from each other as the overlay connections do not reflect the physical connections. This is due to the abstraction offered by the TCP/IP protocol stack at the application layer. Furthermore, by means of cross-layer communication, the overlay network can be matched to the physical network if necessary. This offers important advantages in terms of reduction of the signaling traffic.

Overlay networks allow designers to develop own routing and packet management algorithms on top of the Internet. A similar situation happened in fact with the Internet itself. The Internet was developed as an overlay network on top of the existing telephone network, where long-distance telephone links were used to connect IP routers. Overlay networks operate in a similar way, by using the Internet paths between end-hosts as "links" upon which the overlay routes data, building so a virtual network on top of the network. The result is that overlay networks can be used to deploy new protocols and functionality atop of IP routers without the need to upgrade the routers. New services can be easily developed, with own routing algorithms and policies.

Generally, there are two classes of overlay networks, i.e., routing overlays and storage and lookup overlays. Routing overlays operate on inter-domain IP level and are used to enhance the Border Gateway Protocol (BGP) routing and to provide new functionality or improved service. However, the overlay nodes operate, with respect to each other, as if they were belonging to the same domain on the overlay level. QoS guarantees can be provided as well.

On the other hand, storage and lookup overlays focus on techniques to use the power of large, distributed collections of machines, like in the case of Chord and Akamai. These overlays are actually used as a support for a number of projects on large distributed systems. The distinction between the two classes of overlays has become more and more blurred over the last years.

Strategies for overlay routing describe the process of path computation to provide traffic forwarding with soft QoS guarantees at the application layer. There are three fundamental ways to do routing. These are source routing, flat (or distributed) routing and hierarchical routing. Source routing means that nodes are required to keep global state information and, based on that, a feasible path is computed at every source node. Distributed routing relies on a similar concept but with the difference that path computation is done in a distributed fashion. This may however create problems like distributed state snapshots, deadlock and loop occurrence. There are better versions that use flooding but at the price of large volumes of traffic generated. Finally, hierarchical routing is based on aggregated state maintained at each node. The routing is done in a hierarchical way, i.e., low level routing is done among nodes in the neighborhood of a logical node and high level routing is done among logical nodes. The main problem with hierarchical routing is related to imprecise states.

Notably, overlay routing exploits knowledge of underlying network performance and adapts the end-to-end performance to asymmetry of nodes in terms of, e.g., connectivity, network bandwidth and processing power as well as the lack of structure among them. Overlay routing has the possibility to offer soft QoS provisioning for specific applications while retaining the best-effort Internet model. It can for instance bypass the path selection of BGP to improve performance and fault tolerance.

A specific challenge with overlay routing is related to the presence of high churn rates in P2P networks. The consequence is that the topology is very dynamic, which makes it difficult to provide hard QoS guarantees. Similar situations do exist in wireless ad-hoc networks.

There are two main categories of routing protocols for overlay networks, namely proactive protocols and reactive protocols. Proactive protocols periodically update the routing information, i.e., independent of traffic arrivals. On the other hand, reactive protocols update the routing information on-demand, i.e., only when routes need to be created or adjusted due to changes in routing topology or other conditions (e.g., traffic must be delivered to an unknown destination). Proactive protocols are generally better at providing QoS guarantees for real-time traffic like multimedia. The drawback lies in the traffic volume overhead generated by the protocol itself. Reactive protocols scale better, but they experience higher latency when setting up a new route.

A number of research activities are being carried out worldwide focusing on overlay routing for services like streaming and on-demand. Important research questions are, e.g., on scalability, data search and retrieval, load balancing, churn handling, QoS provisioning with multicast or multipath facilities. It is our ambition to give our contribution to answering these questions.