Internet Content Distribution: Developments and Challenges

Adrian Popescu[†], David Erman[†], Dragos Ilie[†], Doru Constantinescu[†] and Alexandru Popescu[‡]

[†] *Dept. of Telecommunication Systems*
School of Engineering
Blekinge Institute of Technology
371 79 Karlskrona, Sweden

[‡] *Dept. of Computing*
School of Informatics
University of Bradford
Bradford, West Yorkshire BD7 1DP, United Kingdom

***Abstract*— The paper reports on recent developments and challenges focused on multimedia distribution over IP. These are subject for research within the research project ”Routing in Overlay Networks (ROVER)”, recently granted by the EuroNGI Network of Excellence (NoE). Participants in the project are Blekinge Institute of Technology (BTH) in Karlskrona, Sweden, University of Bradford in UK, University of Catalunia in Barcelona, Spain and University of Pisa in Italy.**

The foundation of multimedia distribution is provided by several components, the most important ones being services, content distribution chain and protocols. The fundamental idea is to use the Internet for content acquisition, management and delivery to provide, e.g., Internet Protocol Television (IPTV) infrastructure with Quality of Service (QoS) facilities. Another important goal is to offer the end user the so-called Triple Play, which means grouping together Internet access, TV and telephone services in one subscription on a broadband connection. Other important issues are billing, copyright, encryption and authentication.

The research project is considering the recently advanced IP Multimedia Subsystem (IMS), which is a set of technology standards put forth by the Internet Engineering Task Force (IETF) and two Third Generation Partnership Project groups (3GPP and 3GPP2). IMS offers a wide range of multimedia services over a single IP infrastructure with authentication facilities and, for wireless services, roaming capabilities. Furthermore, the research project is also considering overlay routing as an alternative solution for content distribution.

I. Introduction

Today, the telecommunication industry is undergoing two important developments with implications on future architectural solutions. These are the irreversible move towards IP-based networking and 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) [37]. Taken together, 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. 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 this, the architectural solution must be a unified one, and be 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 being services, content distribution chain, protocols and standards. The basic idea is to use the Internet for content acquisition, creation, management and delivery. An important goal is to offer the end user the so-called Triple Play, which means providing Internet access, TV and telephone services in 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 currently happening in the wide and local area networking is expected to happen in home networking as well. This puts an additional burden on multimedia distribution, which means that wireless access solutions of different types (e.g., WiMAX) must be considered as well. The consequence of adding Triple Play to wireless services is known as Quadruple Play.

It is 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 such as Session Initiation Protocol (SIP) should be taken into consideration. Another important IETF initiative is regarding content distribution issues, which are addressed, e.g., in the IETF WG for Content Distribution Networks (CDN) and Content Distribution Internetworking (CDI). Furthermore, new developments within wireless communications like the IP Multimedia Subsystem (IMS) [10] 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 development and services and cost-effective.

The rest of the paper is as follows. Section II briefly reports on recent developments in CDN as well as on some important challenges related to CDN. Section III is reporting on developments in overlay routing and on important research challenges. Section IV is dedicated to the research project ROVER, and a short presentation of the main research solutions suggested in the project is done. Finally, Section V concludes the paper.

II. 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, as for instance in the case of static or transaction-based Web content, streaming media, real-time video, radio.

There are three distinct categories of content delivery, namely streaming, on-demand and push [23]. The ultimate goal is to optimize the delivery process. The delivery of static, streaming and dynamic content to users is customized in a reliable, secure and scalable manner to allow for more efficient bandwidth management, more intelligent and more flexible content delivery.

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

The fundamental concept is based on distributing content to cache servers located close to end users, thus resulting in better performance, e.g., maximized bandwidth, minimized latency/jitter, improved accessibility. CDNs are composed by multiple Points of Presence (PoP) with clusters (so-called surrogate servers) that maintain copies of (identical) content, thus providing 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 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 [27].

The client-server communication flow is replaced in CDN by two communication flows, namely one between the origin server and the surrogate server and the other 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.

Performance measurements are primarily used to monitor traffic characteristics and gather QoS information about the CDN. Traffic characteristics provide vital clues to the service provider about how the network is being used and they serve as input for network planning (e.g., upcoming hardware and software upgrades). Researchers can build traffic models for various traffic characteristics, which can be used to evaluate existing services or design new ones. For example, at BTH we have performed detailed analysis of BitTorrent and Gnutella traffic that have resulted in parsimonious traffic models suitable for simulation [8], [9], [18], [19].

The QoS information provided by performance measurements can be used to off-load congested portions of the network by re-routing traffic flows and by performing load-balancing [1]. However, this can be quite challenging as in a large network it is not possible to capture a consistent QoS state for the network as a whole. Additionaly, in the case of active measurements the probe rate is a difficult question. Probing the network too often may affect the measured traffic, whereas seldom probing may lead to inaccurate results.

Organizations offering content to geographically distributed clients usually sign a contract with a CDN provider and distribute the content over the CDN by using a specific overlay model. Today, some of the most popular commercial CDN providers are Akamai [2], Nexus [26], Mirror Image Internet [24] and LimeLight Network [22].

In practice, there are several challenges that must be solved in order to offer high-quality distribution at reasonable prices [27], [28]. 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 to improve the performance and what model to use for CDN pricing.

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. Several placement algorithms have been suggested, e.g., Greedy [36], Hot Spot [31] and Tree-Based Replica [21], each of them with own advantages and drawbacks.

Another challenge is the selection of the content to be outsourced 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 [27]. These policies are associated with different advantages and drawbacks. Today however, most of the commercial CDN providers use uncooperative pulling. This is done although non-optimal methods are 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 techniques such as caching, content personalization and data mining can be used to improve the QoS and performance of CDN.

An important parameter to be considered is 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 [16]. It is well known that cost reduction occurs when technology investments allow for delivering services with fewer and cheaper 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 demands for QoS guarantees.

III. 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 [29]. The main advantage of overlay networks is that they offer the possibility to augment the IP routing as well as the 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 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 their own routing and packet management algorithms on top of the Internet. A similar situation happened 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 a virtual network on top of the network. The result is that overlay networks can be used to deploy new protocols and services atop of IP routers without the need to upgrade the routers.

Routing overlays operate on inter-domain IP level and can be 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.

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, e.g., 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 is the need to handle the presence of high churn rates in P2P networks [32]. An important consequence of high churn rates is that the topology is very dynamic, which makes it difficult to provide hard QoS guarantees.

There are two main categories of routing protocols for overlay networks, i.e., proactive protocols and reactive protocols. Proactive protocols periodically update the routing information, independent of traffic arrivals. On the other hand, reactive protocols update the routing information on-demand, 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. Reactive protocols scale better, but they experience higher latency when setting up a new route.

Traffic measurements play an important role in overlay networks, as they are part of overlay routing protocols. Since such protocols do not control the physical links underneath, they typically probe the links to measure parameters such as bandwidth, or latency or packet loss rate. Parameters like average bandwidth, startup time and frame rate are also important for streaming media. Such parameters usually represent desirable or minimal or maximal values that must be obtained in order to, e.g., classify the system response as ”real-time”.

Traffic measurements can be done, e.g., by collecting logs from caches and streaming media servers. They can be also done by deploying software or hardware-based probes throughout the network, especially at the edges of the network. By correlating the information collected by probes with the information collected from cache and server logs, it is possible to determine performance of media delivery and diverse QoS statistics.

As with active traffic measurements, there are important questions that must be answered related to the impact of the measurement probe traffic on network performance, compensation for the effect of measurement traffic, difficulties in mapping large systems, accurate evaluation of the measurement results as well as development of models for adaptive active traffic measurements.

A number of research activities are being carried out worldwide focusing on overlay routing for services such as streaming and on-demand. Important research questions are, e.g., scalability, overlay traffic measurements and modeling, data search and retrieval, load balancing, churn handling, QoS provisioning with multicast or multi-path facilities, congestion and error control in multicast environments [5], [30], [33], [35].

IV. ROVER

The research project ”Routing in Overlay Networks (ROVER)” was granted in 2006 by the EuroNGI Network of Excellence [30]. Participants in the project are Blekinge Institute of Technology in Karlskrona, Sweden, University of Bradford in UK, University of Catalunia in Barcelona, Spain and University of Pisa in Italy.

The main focus in ROVER is on QoS-aware overlay routing in multicast environments, as a way to offer soft QoS provisioning for specific applications while retaining the best-effort Internet model. Main research questions are about overlay traffic measurements and modeling, overlay multicast, QoS provisioning with multicast facilities as well as congestion control in multicast environments.

An important part of our research is on developing a novel class of routing protocols that we are suggesting. For doing this, we use statistics and probability distributions of P2P traffic collected in our measurement studies [8], [9], [18], [19].

The first of the new suggested routing protocols is called the Overlay Routing Protocol (ORP). ORP is a QoS-aware unicast routing protocol, which works in a hybrid fashion, based on ideas used in the wireless ad-hoc routing protocols Associatively-Based Routing (ABR) [34] and the Zone Routing Protocol (ZRP) [14]. The main advantage is that ORP is expected to perform better under churn due to controlled flooding, which reduces traffic overhead in case of rerouting. New routes are setup by a reactive protocol, called the Route Discovery Protocol (RDP), which is based on a flooding algorithm. Furthermore, nodes belonging to an established route exchange routing information among themselves and with their immediate neighbors by using a proactive protocol, called the Route Maintenance Protocol (RMP), which is based on a modified link state algorithm. Proactive protocols for route maintenance offer the advantage that they avoid the latency cost of looking up routes. RMP attempts to repair existing routes or find alternate routes when nodes along a path become unavailable or unable to route according to QoS constraints. In case the RMP fails to repair a broken path, ORP will fallback on RDP.

QoS constraints associated with each route define an optimization problem. To solve this problem, every overlay node maintains measurement information (e.g., bandwidth usage, delay, loss rate) for each traffic flow. The optimization problem can be solved in several ways. For example, the measurement data can be used as input to a Random Neural Network (RNN) that uses this information to continuously adapt the existing routes according to the quality experienced by traffic flows crossing the node. This is done by Reinforcement Learning (RL) [13]. Other methods to solve the optimization problem may be applied as well, e.g., swarm intelligence [6], [7] and genetic algorithms [12]. A comparative study will be carried out on the performance impact of ORP utilizing various optimization algorithms.

Another important part of our research is on additions and improvements of the BitTorrent protocol [4] and piece selection algorithms to make them suitable for streaming media. BitTorrent is a swarming data replication and distribution system, which is based on the game theoretic "tit-for-tat" algorithm. Peers exchange parts of the data, so-called *pieces*, by expressing interest in a given piece from a peer. The suggested protocol extensions aim at improving the process of exchange and allowing for QoS claims and expectations between peers. For instance, a peer could communicate not only interest in a given piece of data, but also add some constraints to this interest, e.g., "I am interested in pieces 3–12, at 25kbps sustained rate". Alternatively, peers will be able to claim to support a specific level of upstream QoS.

Additionally, the standard algorithms for piece selection used in BitTorrent, i.e., Rarest-Piece-First and Random-Piece, are not advantageous for streaming media. This is because the piece reception order is random and streaming applications need a continuous and sequential stream of data to enable non-interrupted playback. We will therefore do research on new selection algorithms (for both peers and pieces) that are suitable for streaming media. We also expect that the *small-world* phenomena often observed in node connectivity [20], i.e., high degree of clustering, will positively impact the stream segment distribution. For instance, well-known stream merging techniques such as patching [17] could be used in this case to patch a received stream more rapidly, and with less server load, than in the case of using a single unicast stream to the original server.

The BitTorrent extensions are part of a research framework on overlay multicast, to handle the last hop distribution, and on a separate set of protocols for the core multicast forwarding and caching.

V. Conclusions

The paper has reported on recent developments and challenges focused on multimedia distribution over IP. These are subject for research within the research project "Routing in Overlay Networks (ROVER)", recently granted by the EuroNGI Network of Excellence.

The main focus in ROVER is on QoS-aware overlay routing in multicast environments, as a way to offer soft QoS provisioning for specific applications while retaining the best-effort Internet model. Main research questions are about overlay traffic measurements and modeling, overlay multicast, QoS provisioning with multicast facilities as well as congestion control in multicast environments.

Planned future work is to develop a dedicated middleware environment, which will be used to develop new protocols for multimedia distribution over IP to offer soft QoS guarantees for specific applications in a multicast environment. We are also planning to develop analytical and simulation models to validate our results.

References

[1] Andersen D., Balakrishnan H., Kaashoek F. and Morris R., *Resilient Overlay Networks*, 18th ACM Symposium on Operating Systems Principles (SOSP), Banff, Alberta, Canada, October 2001

[2] Akamai Technologies, http://www.akamai.com

[3] Biersack E.W., *Where is Multicast Today?*, ACM SIGCOMM Computer Communication Review, Vol 35, No 5, October 2005

[4] *BitTorrent*, http://www.bittorrent.org/index.html

[5] Cui Y., Li B. and Nahrstedt K., *oStream: Asynchronous Streaming Multicast in Application-Layer Overlay Networks*, IEEE Journal on Selected Areas in Communications, Vol 22, No 1, January 2004

[6] Di Caro G., Ducatelle F. and Gambardella L. M., *Anthocnet: An Ant-Based Hybrid Routing Algorithm for Mobile Ad-Hoc Networks*, 8th International Conference on Parallel Problem Solving from Nature, Birmingham, UK, September 2004

[7] Engelbrecht A. P., *Fundamentals of Computational Swarm Intelligence*, John Wiley & Sons, Ltd., 2005

[8] Erman D., *BitTorrent Traffic Measurements and Models*, licentiate thesis, Blekinge Institute of Technology, Karlskrona, October 2005

[9] Erman D., Ilie D. and Popescu A., *BitTorrent Traffic Characteristics*, IEEE International Multi-Conference on Computing in the Global Information Technology (ICCGI'06), Bucharest, Romania, August 2006

[10] Geer D., *Building Converged Networks with IMS Technology*, Computer, IEEE, November 2005

[11] Gelenbe E., Lent R. and Xu Z., *Design and Performance of Cognitive Packet Networks*, Performance Evaluation, No. 46, 2001

[12] Gelenbe E., Gellman M., Lent R., Lei P. and Su P., *Autonomous Smart Routing for Network QoS*, First International Conference on Autonomic Computing, New York, USA, July 2004

[13] Gelenbe E. and Lent R., *Power-Aware Ad-Hoc Cognitive Packet Networks*, Ad-Hoc Networks Journal, Vol. 2, July 2004

[14] Haas, Z. J. and Pearlman M. R., *The Performance of Query Control Schemes for the Zone Routing Protocol*, IEEE/ACM Transactions on Networking, Vol. 9, No. 4, August 2001

[15] Hamra A.A. and Felber P.A., *Design Choices for Content Distribution in P2P Networks*, ACM SIGCOMM Computer Communication Review, Vol 35, No 5, October 2005

[16] Hosanagar K., Krishnan R., Smith M. and Chuang J., *Optimal Pricing of Content Delivery Network Services*, 37th Annual Hawaii International Conference on System Sciences, Big Island, Hawaii, January 2004

[17] Hua K.A., Cai Y and Sheu S., *Patching: a Multicast Technique for True Video-on-Demand Services*, MULTIMEDIA'98: Proceedings of the sixth ACM international conference on Multimedia, Bristol, UK, 1998

[18] Ilie D., *Gnutella Network Traffic: Measurements and Characteristics*, licentiate thesis, Blekinge Institute of Technology, Karlskrona, April 2006

[19] Ilie D., Erman D. and Popescu A., *Transfer Rate Models for Gnutella Signaling Traffic*, IEEE Advanced International Conference on Telecommunications (AICT'06), Guadeloupe, French Caribbean, February 2006

[20] Jin S. and Bestavros A., *Small-World Internet Topologies: Possible Causes and Implications on Scalability of End-System Multicast*, Technical Report BUCS-2002-004, Boston University, 2002

[21] Li B., Golin M.J., Italiano G.F., Deng X. and Sohraby K., *On the Optimal Placement of Web Proxies in the Internet*, IEEE INFOCOM 1999, New York, USA, March 1999

[22] LimeLight Network, http://www.limelightnetworks.com

[23] Asynchronous Layered Coding (ALC) Protocol Instantiation, IETF RFC3450, http://www.ietf.org/rfc/rfc3450.txt

[24] Mirror Image Internet, http://www.mirror-image.com

[25] Neumann C., Roca V. and Walsh R., *Large Scale Content Distribution Protocols*, ACM SIGCOMM Computer Communication Review, Vol 35, No 5, October 2005

[26] Nexus International Broadcasting Association, http//www.nexus.org

[27] Pallis G. and Vakali A., *Insight and Perspectives for Content Delivery Networks*, Communications of the ACM, Vol 49, No 1, January 2006

[28] Plagemann T., Goebel V., Mauthe A., Mathy L., Turletti T. and Urvoy-Keller G., *From Content Distribution Networks to Content Networks - Issues and Challenges*, Computer Communications, Elsevier, Vol 29, 2006

[29] Popescu A., *Routing in Overlay Networks: Developments and Challenges*, IEEE Global Communication Letter, IEEE Communications Magazine, Vol 43, No 8, August 2005

[30] Popescu A., Kouvatsos D., Remondo D. and Giordano S., *ROVER . Routing in Overlay Networks*, EuroNGI project JRA.S.26, 2006

[31] Qiu L., Padmanabhan V.N. and Voelker G.M., *On the Placement of Web Server Replicas*, IEEE INFOCOM 2001, Anchorage, USA, April 2001

[32] Saroiu S., Gummadi P. K. and Gribble S. D., *Measuring and Analyzing the Characteristics of Napster and Gnutella Hosts*, Multimedia Systems, Vol. 9, No. 2, pp. 170-184, August 2003

[33] Shi S.Y. and Turner J.S., *Multicast Routing and Bandwidth Dimensioning in Overlay Networks*, IEEE Journal on Selected Areas in Communications, Vol 20, No 8, October 2002

[34] Toh C-K., *Associativity-Based Routing for Ad-Hoc Mobile Networks*, Wireless Personal Communications, Vol. 4, No. 2, March 1997

[35] Wang B., Sen S., Adler M. and Towsley D., *Optimal Proxy Cache Allocation for Efficient Streaming Media Distribution*, IEEE Transactions on Multimedia, Vol 6, No 2, April 2004

[36] Weisstein E., Greedy Algorithm, http://mathworld.wolfram.com/GreedyAlgorithm.html

[37] WiMAX Forum, http://www.wimaxforum.org