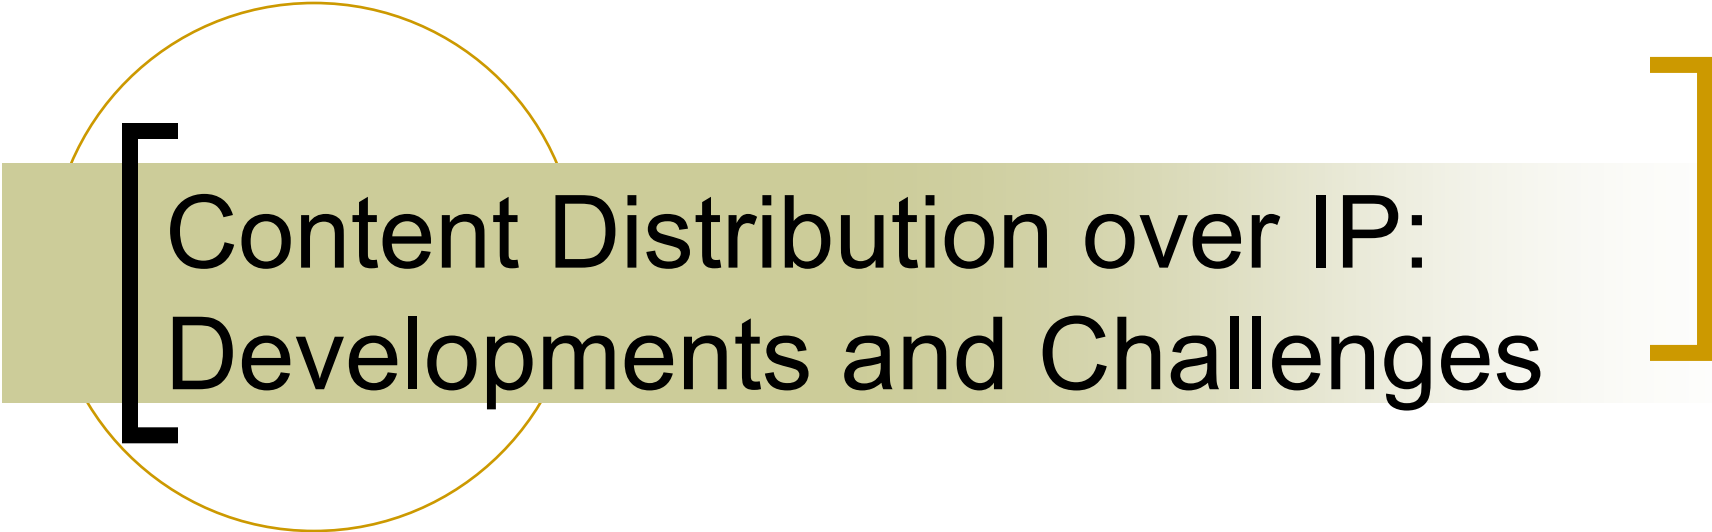




Content Distribution over IP: Developments and Challenges

Adrian Popescu

Blekinge Institute of Technology
Karlskrona, Sweden

October 2008

[Outline]

- Introduction
- Content Distribution Networks, Why?
- Types of CDN
- Cache-Based CDN
- Cache-Based CDN in Practice
- Operational Comparison
- Important Challenges
- ROVER: A .SE Supported Research Project
- Main Scientific Results
- Other Results
- ROVER Architecture
- ROVER Middleware
- Unicast QoS Routing in Overlay Networks
- BitTorrent Media Distribution
- Overlay Multicast Networks
- Conclusions

Introduction


- Currently, three important developments in telecom
 - Irreversible move towards IP- and SIP-based networking
 - Deployment of broadband access, e.g., ADSL2+, FTTH, WLAN
 - Expansion of mobile communication systems, e.g., UMTS, WLAN, WiMAX
- Consequence
 - Appearance of more advanced and more bandwidth-demanding applications, e.g., IPTV, online gaming
- Plethora of requirements
 - Multicast facilities, high bandwidth, low delay/jitter, low packet loss
- More constraints, e.g., billing, authentication, copy right
- Need for unified architectural solution

Content Distribution Networks, Why?

- CDNs are networking solutions where high-layer network intelligence is used to improve the performance in delivering media content over the Internet

- Internet is used for

- Content acquisition
- Content creation
- Content delivery
- Management



CD value chain

- History

- First generation: focused on Web documents
- Second generation: focused on Video on Demand (VoD), audio and video streaming

- Example of offerings:

- Triple/Quadruple Play
- Transaction-based Web content
- Streaming media
- Real-time video, radio

[Types of CDN]

- Unicast CDN
 - Centralized servers host the content
 - Solution good for low-demand content or content that is customized to end user
- Cache-based CDN
 - Content nodes are distributed across the Internet
 - The travel distance of the content is minimized
 - Solution good for high-demand content
 - Two classes available: application-based CDN; application- and network-based CDN
- Peer-to-Peer (P2P) CDN
 - Use of P2P technique
 - More nodes are involved in the delivery process
 - Overlay routing may improve the performance
 - Research done by BTH in the research project ROVER
- Multicast CDN
 - The content is simultaneously sent to end users by adaptively replicating and branching data streams across the IP network
 - The content is delivered only to the requesting end users
 - Solution good for distribution of live streaming and high-demand content
 - Research done by BTH in the research project ROVER


Cache-Based CDN

- Today, cache-based CDN provide some of the best QoS and QoE to large volumes of dispersed end users
- Main concept: they operate by duplicating the content hosted on centralized servers
- Cache servers are distributed globally
- The performance is improved given that, e.g., the content delivery delay and jitter are reduced



Cache-Based CDN in Practice

- Distributing content to cache/surrogate servers located close to end users
- IP routers are used to deliver content to optimal surrogate servers
- Use of overlay routing to obtain better delivery performance
- Client-server communication flow is replaced by two communication flows
- Accounting mechanism to provide logs and other information to the origin servers
- Types of data
 - Encoded/multimedia data
 - Metadata (to allow identifying, finding and managing the encoded data)



Cache-Based CDN in Practice (cont'd)

- Main functional components
 - Content distribution
 - Moving the content from source towards users
 - Request routing
 - Selecting the best location for retrieving the requested content
 - Content routing
 - Delivering the content from the most appropriate place to the client requesting it
 - Content processing
 - Creating and adapting content to user preferences and device capabilities
 - Authorization, authentication and accounting
 - Enabling of monitoring, logging, accounting, and billing of content usage

Cache-Based CDN in Practice (cont'd)

- Approaches to using cache-based CDN services
 - Content owners build and manage own CDN
 - Buy services from specialized CDN owners
 - Also known as application-based CDN
 - Focused on content delivery only, i.e., layers above IP
 - There is no knowledge about the underlying IP network
 - Buy services from CDN owners that offer specialized CDN cache services as a managed service
 - Also known as network-based CDN
 - The CDN owners have all the IP visibility as they own and operate a network infrastructure as well
 - The content nodes use DNS information to approximate the location of an end user

Operational Comparison

Function	Application-layer CDN	Network-Based CDN
Selection of cache server	Based upon the location of DNS for an end user.	Based upon the actual source location of the request.
Rerouting around congestion	Based upon the relative performance of alternative routes.	IP routing is dynamically changed to eliminate congestion within own backbone and peer IP networks.
Content node location	Third-party data centers are used by CDN providers to locate content nodes.	Major routing centers are used as content nodes, where user requests are handled locally.
Load balancing	Content requests are distributed based upon the resources available in cache servers.	Content requests are distributed based upon the resources available in cache servers as well as in access routers, core routers and peering links.
Knowledge of end user	Has no direct knowledge of end user location, it uses DNS information only.	Has direct knowledge of its user and business users by IP address.

Important Challenges

- Where to place the cache servers?
- Which content to outsource and what practice to use for the selected content outsourcing?
- How many CDN architectures do we need, one CDN per application or just one CDN shared by all applications?
- How to exploit P2P technology, overlay routing and caching to improve the performance?
- How to exploit data mining over CDN to improve the performance?
- How to exploit content personalization to improve the performance?
- What model to use for CDN pricing?

ROVER: A .SE Supported Research Project

- ROVER: Routing in OVERlay networks
- .SE supported project 2007-2008
- Focus: media distribution in overlay networks
- Particular focus
 - QoS-aware overlay routing
 - Middleware
 - Mechanisms for media distribution
 - Study of protocols for multicast distribution
- Blekinge Institute of Technology (BTH) team
 - Professor Adrian Popescu
 - TeknDr David Erman
 - TeknDr Doru Constantinescu
 - TeknLic Dragos Ilie
 - PhD student Alex Popescu
 - MSc Karel de Vogeleer


Main Scientific Results

- Three PhD thesis
 - Doru Constantinescu, December 2007, “*Overlay Multicast Networks: Elements, Architectures and Performance*”
 - David Erman, March 2008, “*On BitTorrent Media Distribution*”
 - Dragos Ilie, October 2008, “*On Unicast QoS Routing in Overlay Networks*”
- Two journal papers
 - “Computer Networks” journal, Elsevier, December 2007
 - “Telecommunication Systems” journal, Springer, July 2008
- Twelve conference papers
- Three technical reports

[Other Results]

- Partial implementation of a dedicated middleware
- Framework named Overlay Routing Protocol (ORP) suggested to provide a QoS-aware service on top of IP's best effort service
- Simulation study of ORP
- Modifications and extensions suggested to the BitTorrent (BT) to make it suitable for use in providing a streaming video delivery service
- Simulation study of the suggested BT modifications and extensions
- Comparative simulation study of three representative categories of overlay multicast networks, i.e., Application Layer Multicast Infrastructure (ALMI), Narada and NICE

[ROVER Architecture]



[ROVER Middleware]

- Middleware: software that bridges and abstracts underlying components of similar functionality and exposes this through a common API
- Object-oriented (C++) based API
- Based on the Key-Based Routing (KBR) of the common API framework suggested by the authors of CHORD
- Intended to work on top of both structured and unstructured underlays; compared to this, the initial KBR was suggested to work only on top of a structured underlay
- Quick integration of existing protocol implementations
- Development, evaluation, testing, performance analysis of different protocols and combinations of protocols

Unicast QoS Routing in Overlay Networks

- Particular difficulties
 - Multiple constraints
 - Dynamic environments; presence of churn
 - “Real-time” performance demand
- QoS constraints can be
 - Additive (e.g., for delay), or
 - Multiplicative (e.g., for packet loss), or
 - min-max (e.g., for bandwidth)
- Optimization algorithms
 - Self-Adaptive Multiple Constraints Routing Algorithms (SAMCRA)
 - The Simplex Method
 - Gradient Projection Method
 - Conjugate Gradient Method
 - Particle Swarm Optimization

Unicast QoS Routing in Overlay Networks (cont'd)

- **ROUTING:** the process of selecting paths in a network such that they satisfy a set of simultaneous QoS constraints
 - Routing algorithms: given a network topology, find the desired paths
 - Routing protocols: ensure that all nodes have “accurate” topology information
- Example: *“Find a path from node A to node B with a minimum of 1Mbps capacity, such that the delay does not exceed 100ms and the packet loss probability is no higher than 0.01%”*
- Types of path QoS metrics (example: path $i \rightarrow j \rightarrow k \rightarrow \dots \rightarrow l \rightarrow m$)
 - Additive (delay, jitter): $d(i,j) + d(j,k) + \dots + d(l,m)$
 - Multiplicative (packet loss): $1 - (1 - p(i,j)) \times (1 - p(j,k)) \times \dots \times (1 - p(l,m))$
 - Min-max (bandwidth): $\min(c(i,j), c(j,k), \dots, c(l,m))$

Unicast QoS Routing in Overlay Networks (cont'd)

- Research has been done on
 - Finding paths suitable for transporting multimedia flows
 - Path selection is done such as to satisfy a set of simultaneous QoS constraints
 - Reacting to path failures by reallocating flows to backup paths
 - Implementing this functionality in an overlay network spawned by end-nodes, without changing existing Internet infrastructures

Unicast QoS Routing in Overlay Networks (cont'd)

- Study has been done on
 - Flow allocation problems and optimization algorithms
 - Gnutella measurements and characteristics modeling
 - Overlay Routing Protocol (ORP) framework
 - Route Discovery Protocol (RDP): finds QoS-constrained paths by selective forwarding
 - Route Maintenance Protocol (RMP): handles churn by reallocating flows to backup paths
 - Different performance metrics have been evaluated, e.g.,
 - Call blocking ratio, bandwidth utilization, bandwidth overhead, path stretch (RDP)
 - Path failure ratio, restored paths ratio, bandwidth utilization, bandwidth overhead (RMP)
- The experiments have shown that RDP and RMP are viable alternative to provide a QoS-aware service at the application layer
- The cost has been observed to be maximum 1.5% of the residual network capacity
- Future work regards implementation of RDP and RMP and PlanetLab tests

[BitTorrent Media Distribution]

- Research has been done on BitTorrent (BT) to develop mechanisms to adapt BT for near-zero delay Video on Demand (VoD)
- Contributions
 - Overview of distribution mechanisms for streaming media
 - Comprehensive set of BT traffic models at session and message level
 - Modifications and extensions to core BT algorithms

BitTorrent Media Distribution (cont'd)

- Streaming with BT
 - For content replication
 - Download full pieces quickly
 - Download most rare pieces
 - For streaming
 - Maintain ordering of pieces to be consumed
 - Keep playback deadlines
- BT extensions for streaming: main ideas
 - Decrease content server load
 - Leverage existing BT peers
 - Decrease traffic variability by smoothing
- Modifications needed
 - Piece selection
 - Peer selection
 - Provision of QoS guarantees in data transfer

BitTorrent Media Distribution (cont'd)

- Simulation results have shown the feasibility of this solution
- Research continued in another .SE project supported 2008/2009 on “Mechanisms for Media Distribution (BiTES)”

[Overlay Multicast Networks]

■ Advantages

- The network only needs to provide the basic stateless, unicast, "best-effort" delivery
- No special infrastructure is needed; ease of deployment
- Services are provided over heterogeneous networks
- No need for special operating system support; higher flexibility, soft QoS possible


■ Drawbacks

- Less efficient than IP-level multicast
- Existing applications may need to be modified
- Each application builds its own overlay network
- Existing IP multicast-based protocols cannot be used

Overlay Multicast Networks (cont'd)

- Comprehensive simulation study has been done on the performance of three fundamental Application Layer Multicast (ALM) solutions
 - Tree-first: Application Level Multicast Infrastructure (ALMI)
 - Mesh-first: Narada – an implementation of End System Multicast (ESM)
 - Hierarchical: NICE is the Cooperative Internet Environment (NICE)
- **Myns** simulation engine has been used; GT-ITM topologies integrated into the simulation engine
- Churn models have been implemented
- Performance metrics used: link stress; link stretch; resource usage; control overhead; efficiency.
- Protocol performance in terms of
 - Protocol scalability
 - Protocol dynamics
 - QoS provisioning capability
- New results have been obtained on
 - Protocol scalability with regard to the type of application used
 - Overlay multicast group management
 - Resource usage
 - Robustness to churn
 - QoS provisioning capability

Overlay Multicast Networks (cont'd)



Simulation model

Conclusions

- Very fascinating and complex research!
- This research has resulted in
 - Three PhD thesis
 - Two journal publications
 - Twelve conference publications
 - Three technical reports
 - The .SE research project “Mechanisms for Media Distribution (BiTES)”
 - BTH is participating in the Eureka project “Mobicome” with 13 partners from Norway, Sweden and Spain
 - BTH is participating in the EU FP7 STREP project “Perimeter” with 15 partners from different European countries
- Need for move towards real-live deployment, e.g., PlanetLab in a first phase
- Need for participation in the standardization efforts

THANK YOU!

