Architectural Trade-offs for Video Transport Networks

Didier Colle, Goutam Das, Mario Pickavet and Piet Demeester

IBBT – Ghent University, Department of Information Technology (INTEC)
Gaston Crommenlaan 8 bus 201, B-9050 Gent, Belgium
{didier.colle, goutam.das, mario.pickavet, piet.demeester}@intec.UGent.be

Abstract: several architectural aspects and options are discussed for producing video based services: storage and transcoder placement, routing and (staggered) multi-casting, buffering/smoothing and retiming streams, separate handling of I-frames from B-/P-frames and applying optical flow switching.
©2010 Optical Society of America
OCIS codes: (060.4250) Networks; (060.1155) All-optical networks

1.   Introduction

It is a well-known fact that the Internet traffic has been growing very fast in the past. Whereas a decade ago, the Internet traffic was doubling every year (+100% p.a.), it is still expected to grow at a very fast pace of doubling every two years (+50% a +60% p.a.). Whereas peer-to-peer (P2P) file sharing was the dominant traffic boosting the Internet traffic growth, video traffic will take over this role in the coming years.

Already today, video traffic is an important fraction of the total Internet traffic. The MSN Video Service handled around 500 million requests in the period April-December 2006 while the repository contained around 59000 video files [1]. This boils down to 20 requests per second or a request each 50 ms. In January 2009, 14.8 billion YouTube views were recorded [2], corresponding to more than 5500 requests per second or a request each 180 µs.

Not only online viewing video files over the Internet via a PC or laptop is popular. Also Video-on-Demand (VoD) services to TV viewers at a decent quality are becoming rapidly popular. For example, the cable operator Telenet in Flanders (6 million inhabitants) in Belgium launched a few years ago the Digital-TV service to which a third of the TV services customers had subscribed in 2008: these customers were requesting on-demand 20 million videos during 2008 [3] equaling to a request each 1.5 sec.

In addition to the growing popularity of all kind of video services, also the bitrate per video stream keeps increasing, as consumers are expecting higher and higher resolutions at an increasing quality. Nowadays, High-Definition (HD) TV is already replacing the Standard Definition (SD) TV, requiring 10-20 Mbps instead of 3-4 Mbps. Evolutions to 4K, 8K and even 3D television can be expected in the coming decade(s). Despite that codecs are continuously improved (coding efficiency of doubling each decade or only increasing ~7% per year [4]), they will not be able to keep up with the ever increasing resolution/quality requirements of new video standards.

Summarizing, the expectation is that video traffic will dominate the Internet traffic in the future and thus it becomes urgent to optimize the network infrastructure according to this evolution. More concretely, according to forecasts by Cisco [5] global IP traffic will be growing from 10 Exabyte (EB) per month in 2008 to 56 EB per month in 2013: the majority of the traffic being consumer traffic, growing from 7 EB per month to 41 EB per month. Whereas file sharing was the largest fraction in 2008 (3.4 EB per month) it will only grow to 10 EB per month while it will be dominated by all kinds of video traffic reaching at least 26 EB (of which only 8.4 EB is confined to a single service provider: e.g., the service provider VoD or IPTV service) per month in 2013.

2.   Optimizing video transmission and delivery networks

An important distinction to make in video traffic is real-time versus non-real-time viewing of the video traffic.

Video conferencing and TV broadcasting are exemplary services of real-time video viewing. In such context, the video source (e.g., camera) is directly streaming the video content to the display. Instead of setting up a dedicated stream per receiver (display), the stream may be multi-/broad-casted inside the network. As in deployed network systems often no multicast features are implemented, often multicast is considered at the application level [6], [7].

Of course, services like video conferencing will not disappear (on the contrary), but it is expected that more and more viewing behavior will shift to non-real-time viewing of content (time-shifted TV is just one example). Therefore, this paper is mainly focusing on this kind of video services.

As illustrated in Fig. 1, several steps may be crossed before video content reaches the receiver on which it is displayed. Video cameras and more generally video production systems ingest video content on content servers inside the network. In the simplest case, the video content is fetched from these content servers each time the customers want to view a video file. Either the video content is streamed in real-time (signaling protocols often used are Real-Time Streaming Protocol (RTSP) or the Session Initiation Protocol (SIP)) or is downloaded like a regular webpage (e.g., over HTTP). The challenge from a network perspective in the former case is that the network should be capable to guarantee stringent QoS (jitter and delay, packet loss, …) typically required for video streaming, while in the latter case the network should be able to provide enough capacity (in order to facilitate sufficient interactivity for fast-forwarding or for unavoidable downloading of content that is not viewed by the customer itself).

In order to leverage capacity stress on the core network, popular content may be cached (i.e., temporary stored) on rather small caching servers: it appears that only a minor fraction of all available video content is accessed very frequently [1]. Another way to reduce the capacity stress is not to establish a dedicated connection/stream for each request, but to apply staggered multicasting for popular content: the same content is streamed and multi-/broad-casted several times but slightly shifted in time (e.g., restarting the video content every 5 minutes) in order to allow the customer to tune in on the stream he desires. A disadvantage of this approach is that it would be applied to popular content and thus implying that receipt of this popular content can be slightly delayed. A slightly different approach called Video-In-Network (VIN) [8] is to have the video content circling around in the network, thus abusing the network itself as storage system. Although high-end video servers are rather expensive, it is pretty questionable this solution can ever become cost-efficient: storing half an hour SDTV at 3.75 Mbps would require multiple parallel 100 Gbps line-systems with a reach of a couple of 1000 kilometers (otherwise, profiting of buffering delays would be comparable to a RAM in the content servers).

Finally, approaches based on any kind of multi-/broad-casting may suffer from the problem that the content should be displayed on different kind of receiving devices (e.g., large TVs versus mobile phones). Therefore, some nodes in the network may transcode a video stream in such a way that it (down-)converts the stream to one that is compatible with the displaying device and/or getting through the access network bottleneck. Transcoders require significant amounts of processing power, thus network cost has to be traded against introducing processing power closer to the costumer for transcoding purposes.

The above discussion can be generalized in following optimization problem: where to install content and caching servers and transcoders, how to handle requests based on the popularity of the requested content and how to route the content between these elements across the network.
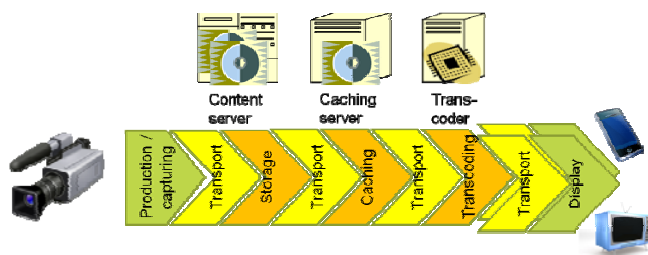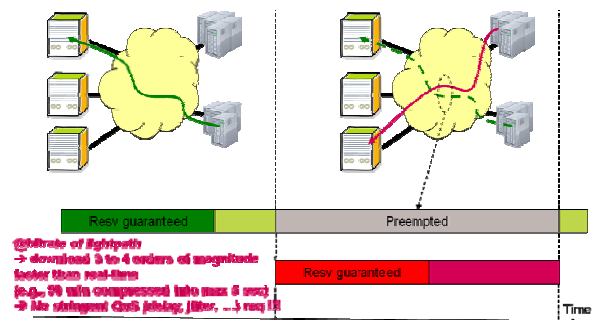

Fig. 1. Video production and delivery chain

Fig. 2. Optical flow SwItching for faster than Real-time vIdeo diStribution (OSIRIS) concept.

3. Video encoding and the relation to video transmission

Video streams are often very bursty. This is due to the fact that most video frames are P- or B-frames, which represent the difference with previous/next frames, whereas only a small number of independent I-frames are present.

A solution to overcome the burstiness is to buffer/smooth the traffic at the ingress of the network and to reconstruct the actual bursty stream at the egress of the network: such an approach would require significant amounts of buffering capacity and introduce according delays. Another approach could be to separate transmission of I-frames and P-/B-frames. A slightly different approach would be to have a stream of only P-/B-frames and only to deliver large I-frames on-demand: this would of course require that such functionality is enabled in the codec at the source of the video stream.

4. Role of optical networks in video delivery: the OSIRIS concept

It is obvious that optical fiber communication is crucial in cost-efficiently deploying the bandwidth hungry video services. However, also the optical switching technologies like the MEMS-technology might be interesting.

The idea of the Optical flow SwItching for faster than Real-time vIdeo diStribution (OSIRIS) concept illustrated in Fig. 2 is to establish a lightpath per video file to be transferred and to push it as fast as possible through that ligthpath, so that when another request arrives the receiver has enough content received so that the lightpath could be temporary pre-empted. In this way, one can avoid that a very huge amount of rather small packets need to be processed identically by L2/L3 devices.

Latching-capable MEMS-switches require only a small amount of energy for reconfiguring the switch, resulting in several orders of magnitude improvement in energy efficiency compared to a L2/L3 switching network. Given that half an hour of SDTV would require several tens of ms at 100 Gbps would make the MEMS-technology with its switching time in the order of ms sufficiently fast.

A problem of the OSIRIS concept might be that it stresses in terms of access speed too much the storage/memory at the head and tail of the established lightpath. Nevertheless, given the fact that a single file needs to fill up the bandwidth of the lightpath, the access speed of the storage/memory at the edge of network could be dimensioned according to the peak rate rather than the sustainable rate as is typically done under the condition the content can be stored as contiguous blocks. In other words, the OSIRIS concept would facilitate adopting storage devices featuring immense peak/sustainable rate ratios.

5. Conclusions

Video traffic will drive the Internet traffic growth and dominate the overall traffic, calling for video optimized network solutions. The paper shows that routing, choosing between uni-/multi-/broad-casting, placement of content and caching servers and transcoders need to be optimized jointly. Several options to transport video streams across the network taking into account the actual coding of video in large independent I-frames and rather small relative B-/P-frames are also discussed. Finally, the OSIRIS concept based on optical flow switching for enabling faster than real-time video distribution is presented as a promising novel network concept.

6. References

[1]   Cheng Huang, Jin Li, Keith W. Ross, "Can Internet Video-on-Demand be Profitable?", SIGCOMM2007
[2]   "YouTube reaches 100 million U.S. viewers", http://youtubereport2009.com/youtube-reache-100-million-us-viewers/
[3]   "Substantial increase for Telenet Video On Demand service", http://hugin.info/136600/R/1276641/284338.pdf
[4]   M. Ghanbari et al, "Future Performance of Video Codecs" (SES2006-7-13), November 2006, http://privatewww.essex.ac.uk/~fleum/FinalOfcomReport.pdf
[5]   "Cisco Visual Networking Index:Forecast and Methodology, 2008–2013", http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf
[6]   C. K. Yeo et al, "A survey of application level multicast techniques" in Computer Communications 27 (2004) 1547–1568.
[7]   D. Cheng et al, "A Solution to Select Multicast Service Nodes of Hierarchical Overlay Multicast Tree Based on Immune Evolution" in IEEE Fourth International Conference on Natural Computation, 2008.
[8]   H. Q. Guo, L. H. Ngoh, W. C. Wong, J. G. Tan, "Comparison of in-network versus Staggered Multicast video distribution models", Multimedia Tools and Applications, March 20086