SEACAST: A PROTOCOL FOR PEER-TO-PEER VIDEO STREAMING SUPPORTING MULTIPLE DESCRIPTION CODING

Simone Zezza, Enrico Magli*, Gabriella Olmo*, Marco Grangetto†*

* Dip. di Elettronica
Politecnico di Torino
name.surname@polito.it

†Dip. di Informatica
Università degli Studi di Torino
marco.grangetto@di.unito.it

ABSTRACT

SEACAST is a peer-to-peer live streaming protocol developed at Politecnico di Torino, which aims at improving current systems in two key areas. The first is the use of full-fledged flow control using RTP/UDP and session signaling. The second is the use of multiple description coding to handle error resilience and user heterogeneity. In this paper we overview SEACAST, highlighting its main innovations, and providing a short summary of performance evaluation over a local testbed at Politecnico di Torino. The results show a definite performance improvement with respect to existing systems, and point out the usefulness of multiple description coding in the peer-to-peer context.

1. INTRODUCTION

This paper introduces SEACAST, a new peer-to-peer (P2P) live streaming protocol developed in the framework of the European Project "SEA - seamless content delivery". SEACAST represents an evolution of the VidTorrent protocol developed by Viral Communication [1, 2]. Similarly to other P2P streaming protocols, VidTorrent's design is still largely based on file-sharing systems. The objective of SEACAST is to overcome this limitation, adopting techniques and protocols optimized for H.264/AVC video streaming. In particular, SEACAST introduces two main innovations, namely flow control using a proper protocol stack, and application-layer error control by means of state-of-the-art multiple description coding (MDC) techniques.

As for flow control, in P2P video streaming the overlay network should guarantee a constant flow of data and a low start-up latency, while coping with a significant packet loss rate. While many P2P systems employ TCP for video transport, SEACAST breaks this legacy of file-sharing systems, and uses the RTP, UDP, and RTSP/SDP to set up sessions and enable flow control.

The research leading to these results has received funding from the European Community's Seventh Framework Programme under grant agreement 214063 (SEA - seamless content delivery).

As for error control, the presence of prediction loops in the H.264/AVC encoder induces error propagation, and retransmission mechanisms cannot be afforded due to the constraints of real-time delivery. In SEACAST, MDC is used to cope with this vulnerability. MDC is particularly suitable for multiple tree overlay network topologies. In fact, if each description is delivered over a different tree, a suitable level of diversity is achieved, which necessary to exploit MDC. To overcome the performance loss of MDC with respect to the corresponding single-description coding, in SEACAST we employ the highly efficient state-of-the-art DIVA MDC algorithm [3], in connection with a multiple-tree topology. DIVA is tailored to H.264/AVC, and exploits the concept of redundant slices to perform MDC in a syntax-compliant way. The objective is to assess the advantages provided by SEACAST, and specifically to evaluate the improvement yielded by MDC over a multiple-tree topology on the service continuity, enhancing the final user experience.

2. THE MDC ALGORITHM

In MDC, several non hierarchical, independently decodable representations of the same data (descriptions) are generated, yielding mutually refinable video streams. Such descriptions are suitable for transmission over independent paths, and usually have the same importance in terms of the quality of the recovered data. The DIVA algorithm [3] aims at creating two balanced descriptions, each of which is an H.264/AVC compliant bitstream. To this purpose, it uses the concept of primary and redundant slices [4], creating high-quality primary slices and low-quality redundant slices. This is achieved performed the encoding with two different QPs. Primary and redundant slices are interleaved to form two balanced descriptions. At the decoder, if both descriptions of a given slice are received, only the primary representations are retained and decoded. If a description is lost, the received one is a compliant H.264/AVC bitstream containing primary and redundant representations of alternate slices, which can still be decoded yielding inferior quality.

3. THE SEACAST PROTOCOL

SEACAST is an evolution of the VidTorrent protocol. VidTorrent builds an adaptive overlay for real-time streaming, which aims at minimizing the latency, and is designed for high packet loss rate, and dynamic and heterogenous populations. Data distribution uses a dynamic multiple-tree topology. Trees are built based on bandwidth availability and proximity-aware clustering. The protocol is distributed, as each peer has only local knowledge of the overlay. In VidTorrent, the input stream is split into substreams, which are distributed over independent trees. Typically, each substream contains every other frame of the video sequence. However, the substreams cannot be independently decoded, leading to significant error propagation in the event of packet losses.

SEACAST improves on VidTorrent in a number of ways, as described below.

- VidTorrent uses the HTTP/TCP protocol stack for data streaming. However, this is not suitable for real time video application. In SEACAST, video streaming has been modified so as to use RTP/UDP data flows. This also allows to use MDC for error resilience. In perspective, MDC could also be used to serve heterogeneous users, each user requesting a different number of description based on their bandwidth availability and quality requirements.
- RTSP/SDP signalling has been implemented to describe streaming sessions.
- DIVA MDC has been integrated with SEACAST, and replaces the trivial odd/even frame data splitting present in VidTorrent.
- VidTorrent has major problems to scale to a large number of nodes. Even though these problems are not completely solved, at present SEACAST is able to scale to a sufficient number of nodes (about 50) to enable validation on a local test-bed; full functionality is expected to be available shortly.

4. EXPERIMENTAL RESULTS AND CONCLUSIONS

We have evaluated the performance of SEACAST on a local testbed, consisting of several personal computers running Linux or Linux virtual machines. The testbed can simulate up to about 50 peers in real-time. The peers are connected through the Ethernet of Politecnico di Torino, and each peer has a software module allowing to modify the available incoming and outgoing bandwidth and delay on-the-fly. Each node runs the SEACAST software. Upon signing in to the server, a node is redirected to the SEACAST broker that connects her to the overlay. The broker provides a list of peers from which the user can receive the desired video. In MDC mode, the SEACAST server manages a multiple tree of order 2, and sends either description on a separate tree. Each description has a bit-rate of 512 kbit/s.

In order to assess the effectiveness of MDC, we have emulated the following scenario. Fifty peers connect to the server at random times, requesting the same video stream. Congestion are simulated on 10 randomly selected peers, such that the available bandwidth decreases from 1 Mbit/s to 400 kbit/s. The server periodically makes a rough estimation of the available bandwidth using `http` packets. When the server detects the reduced bandwidth, it disconnects the peer from one of the trees, as she does not have enough bandwidth to receive the full video. As a consequence, this peer only receives one degraded description, and all her children on the remaining tree are negatively affected.

We have measured the number of descriptions received by all peers over a time interval of 45 minutes, through analysis of log files of each peer. It has been found that, using MDC, no peer suffered a service interruption due to congestion, as all of them were able to receive at least one description. Conversely, the original VidTorrent configuration suffered from severe quality degradation upon congestion of a link, as the effect would be a video freeze.

These preliminary results show that MDC is a useful tool to improve the continuity of service in P2P live video streaming systems. Future developments will be in the direction of generating more than two descriptions/trees. Moreover, it is foreseen to provide support for scalable video coding by modifying the policies for tree generation, so as to achieve trees with differentiated reliability. Finally, SEACAST will also be tested on large-scale P2P scenarios, e.g., using Planetlab.

5. REFERENCES

[1] "http://viral.media.mit.edu/index.php?page=vidtorrent".

[2] Viral Communications, "http://viral.media.mit.edu".

[3] T. Tillo, M. Grangetto, and G. Olmo, "Redundant slice optimal allocation for H.264 multiple description coding," *IEEE Trans. on Circ. and Syst. for Video Tech.*, vol. 18, no. 1, pp. 59 – 70, Jan 2008.

[4] T. Wiegand, G.J. Sullivan, G. Bjntegaard, and A. Luthra, "Overview of the H.264/AVC video coding standard," *IEEE Trans. on Circ. and Syst. for Video Tech.*, vol. 13, no. 7, pp. 560–576, Jul 2003.