

Distributed Optimization of Media Flows in Peer-to-Peer Overlay Networks

Antonios Argyriou
Philips Research
Eindhoven, The Netherlands
Email: antonios.argyriou@philips.com

Jacob Chakareski
Ecole Polytechnique Fédérale de Lausanne (EPFL)
Lausanne, Switzerland
Email: jacov.cakareski@epfl.ch

Abstract—We consider the problem of rate-distortion (RD) optimized media streaming in unstructured peer-to-peer (P2P) overlay networks. We formulate the aforementioned problem as a distributed rate allocation problem, and we solve it by applying classical decomposition techniques so that the network-wide utility of the media distortion is minimized. Information exchange between the peers is employed to ensure updates on the price of the locally calculated rate allocation. Media packets are also piggybacked with RD preambles that contain information regarding their impact on the decoder distortion and their size. The benefit of the aforementioned approach is that peers can convert the calculated optimal rate allocation into simple forwarding or dropping actions allowing thus a lightweight implementation. Our simulation results indicate that significant quality benefits can be achieved when the precise RD characteristics of a media description are taken into account by the streaming algorithm.

Index Terms—Video streaming, peer-to-peer network, distributed rate allocation, rate-distortion model.

I. INTRODUCTION

Peer-to-peer (P2P) networks have emerged as an alternative solution to IP multicast for point-to-multipoint media distribution [1], [2]. A P2P network is basically an overlay network that consists of unicast sessions between the cooperating peers. The main task of the peers is to act as proxies that cache a portion of the received packets which are then forwarded to other peers in the network. One of the main benefits of this delivery model is that it offers a scalable way of distributing on-demand or live video over a large number of receivers since the capacity of the system grows with the number of peers. Large-scale on-demand and live P2P multimedia streaming are used with great success [3]. The performance characteristics of these systems, depend primarily on the overlay construction and maintenance algorithms [4], [5].

Probably the most important problem with P2P networks is that the structure of the overlay is affected by the dynamic behavior of the participating peers. Therefore, if the overlay is organized as a tree, bandwidth fluctuations and peer failures that reside close to the root of the tree may cause buffer underflow at a large group of downstream nodes. To alleviate these problems, mesh-based overlay protocols enable data dissemination between peers with no explicit support from a regular overlay structure [3]. With this approach, a peer randomly selects a subset of target peers to push recently received media packets to them while simultaneously receiving




Fig. 1. Simple scenario that indicates media distribution through an overlay. The overlay node P4 must deliver the packets that maximize the quality both at clients C1 and C2.

segments pushed packets from other peers. The *gossiping* process that precedes the actual data delivery is responsible for the highly robust data distribution. Nevertheless, it is not straightforward to apply this delivery model for on-demand streaming, since it may fail to achieve a timely delivery. Another disadvantage is that by allowing gossiping between neighboring peers, excessive data duplication could occur which can be an important overhead for high-bandwidth video streaming applications. Despite the wealth of research on P2P media streaming, most of the existing works have focused on the performance of different application-layer overlay delivery models [2]. Although very important, the communication model and the overlay maintenance algorithms have been developed and studied independently from the actual nature of the content that is being disseminated. More specifically, there is a lack of a generic framework that considers the impact of individual media packets on the quality of P2P media streaming.

In this paper, we take one step in this direction by adopting a distributed media flow optimization approach that employs rate-distortion (RD) optimized packet scheduling and streaming that is independent of the overlay structure. RD-optimized streaming algorithms have been successfully employed for end-to-end applications [7], and therefore they provide another way of studying media distribution under a more generalized communication model like a P2P overlay network. In the majority of P2P media streaming works, the level of abstraction at which media flows are modeled only

considers the QoS requirements like maximum acceptable delay, minimum bandwidth, and minimum packet loss rate. However, the precise impact that particular media packets have on the quality of a decoded media stream has not been taken into account. Media packets have different importance with respect to their contribution to the rate-distortion (RD) performance of a decoded media stream [7]. Only recently there was some work towards this direction. For example in [6] the authors considered also RD-optimized streaming in a P2P network, but they focused on progressive and multiple description coding. We believe that taking into account the RD characteristics of media packets is an important concept that allows the evaluation of different overlay topologies on the quality of media streaming at a finer granularity.

II. SYSTEM MODEL

To avoid limiting our scheme in specific topologies, we assume a mesh-based P2P delivery model where participating peers form a randomly connected and directed mesh (i.e., unstructured overlay). The connection between two peers is uni-directional which means that data is delivered from a parent to a child peer. Each peer in the overlay has multiple parents and multiple children. We assume that a peer can obtain a list of the currently active peers from a central server. Specifically, this bootstrapping server maintains a list of all participating peers and provides a subset of participants to each new peer. Each of the participating peers forwards a subset of the media flows that are delivered through the same overlay. Such an overlay topology can be modeled as a directed acyclic graph $G = \{N, A\}$ where N is the set of all overlay nodes while A is the set of directed links connecting peers between them. Let us also define as M the number of media flows currently being delivered by the overlay. Also, let A_m denote the set of links between peers that are used for delivering media flow m . Fig. 1 presents a simple topology with two media flows being delivered through the overlay.

We consider media flows at the packet-level since different media packets have different impact on the video quality if they are not decoded. Let the index of a media packet from media stream m be denoted as k_m . Similar to related works, e.g. [7], the RD information associated with this packet consists of its size $R(k_m)$ and the importance of the packet for the distortion of the reconstructed video stream that is denoted as $\Delta D(k_m)$. In practice $\Delta D(k_m)$ is the total increase in the mean square error (MSE) distortion that will affect the video stream if the packet is not delivered to the receiver by its prescribed deadline [7]. In the next section, we proceed with the definition of the global rate allocation problem.

III. PROBLEM FORMULATION

The objective of the proposed approach is to calculate the optimal bandwidth allocation for each media flow that is forwarded by a peer [8]. This is a distributed rate allocation problem that can be formulated as a constrained optimization of a utility function. Before proceeding with the formal definition of the aforementioned problem, we introduce the

necessary notation first. Let $D_m^{(i,j)}(r_m^{(i,j)})$ denote the distortion associated with flow $m \in \mu_i$ when it is sent over the link between peers i and j at a rate $r_m^{(i,j)}$. Let also μ_i denote the group of incoming media flows at peer i . Then, the overall distortion of all flows being delivered over the link (i, j) can be expressed as

$$D^{(i,j)} = \sum_{m \in \mu_i} \gamma_m D_m^{(i,j)}(r_m^{(i,j)}), \quad (1)$$

where γ_m is an importance factor that peer i can assign to media flow m when performing rate allocation over the outgoing link (i, j) . This weight factor expresses the importance of flow m in terms of its delivery requirements. It indirectly affects the rate that is allocated to that particular flow. In addition, let

$$r^{(i,j)} = \sum_{m \in \mu_i} r^{(m)}(i, j) \quad (2)$$

denote the overall rate of the flows sent over link (i, j) , where $r_m^{(i,j)}$ is the rate assigned to flow $m \in \mu_i$. Using the previous expression for the total distortion that is introduced at a particular link between two peers, we can express the average media distortion for all peers and all the flows that are transported in the overlay network as follows:

$$D_N = \sum_{(i,j) \in A} D^{(i,j)} \quad (3)$$

We are interested in minimizing the network distortion D_N such that the available bandwidth $R^{(i,j)}$ on each link $(i, j) \in A$ is not exceeded. Hence, the optimization problem under consideration can be formally written as follows:

$$\begin{aligned} \min D_N \\ \text{s.t. } r^{(i,j)} \leq R^{(i,j)}, \quad \forall (i, j) \in A \end{aligned} \quad (4)$$

In the next section we focus on deriving a distributed solution to the previous problem.

IV. DISTRIBUTED SOLUTION

A. Global Rate Allocation

Now the goal is to allow each peer to solve the optimization problem with minimal coordination and message passing [9], [10]. Since the RD curve is concave and twice differentiable we can directly apply Lagrange duality for solving the previous constrained optimization problem. Furthermore, the utility function is separable and a distributed solution to this problem can be derived after the dual decomposition. We can apply Lagrange duality to the constraint in (4), and produce the partial Lagrangian as follows:

$$L = D_N + \sum_{(i,j) \in A} \Lambda_{(i,j)} r^{(i,j)} \quad (5)$$

In this equation $\Lambda_{(i,j)} > 0$ is the Lagrange multiplier for link (i, j) at peer i . The dual function is then defined as

$$g(\Lambda) = \max_{\Lambda > 0} L(\Lambda), \quad (6)$$

where Λ is the vector of the Lagrange multipliers at every peer. Therefore, the dual problem is:

$$\begin{aligned} \min g(\Lambda) \\ \text{subject to } \Lambda \geq 0 \end{aligned} \quad (7)$$

Now the Lagrange multiplier expresses the "price" of each selected rate allocation. It is known that if Λ^* is the optimal solution for the dual problem, then $r^{*(i,j)}(\Lambda^*)$ is the optimal solution to the primal problem defined in (4). Furthermore, this Lagrange multiplier decomposes the primal problem into individual streaming rate allocation problems at every peer that can be optimized individually. In particular, each peer i must solve for the optimal rate allocation:

$$r^{*(i,j)} = \arg \min \left\{ D^{(i,j)} + \Lambda_{(i,j)} r^{(i,j)} \right\} \quad (8)$$

The convergence of the dual algorithm to optimal solution has been proven for distributed flow control problems with a convex utility function [11]. For calculating the Lagrange multiplier $\Lambda_{(i,j)}$ at every peer i , we employ a gradient method [12] that updates $\Lambda_{(i,j)}$ according to

$$\Lambda_{(i,j)} = \max \left\{ 0, \Lambda_{(i,j)} - \beta \left(R^{(i,j)} - r^{*(i,j)} \right) \right\} \quad (9)$$

every time after rate allocation is performed.

B. Multi-Flow Rate Allocation at a Single Peer

At the next level of the optimization, the optimal rate allocation is calculated for the flows that are sharing a common outbound link at a peer. This is done based on the importance of the forwarded media flows and reflected through the gradient factor $\lambda_m = \delta D_m(R) / \delta R$. This factor describes the trade-off between data rate and signal distortion for a media flow m . In particular, on the operational rate-distortion curve $D_m(R)$ for flow m we select the point where it holds $\gamma_m \delta D_m(R) / \delta R \approx \Lambda_{(i,j)}$. The rate value associated with this point on the curve is the optimal value that should be assigned by peer i to the incoming flow j on the outgoing link (i,j) . Let this optimal value be denoted as $r_m^{*(i,j)}$. Then, the overall forwarding rate for the peer on the link introduced in (8) is computed as

$$\sum_{m \in \mu_i} r_m^{*(i,j)}. \quad (10)$$

Since a media flow is already assigned its outgoing rate $r_m^{*(i,j)}$, a peer can easily allocate this rate to every individual packet comprising that flow without any further optimization steps. Recall from Section II that information on the rate-distortion characterization is piggybacked with each individual packet. Since each packet k is labeled with $\Delta D(k_m)$ and $\Delta R(k_m)$, a peer can easily calculate the per-packet gradient $\Delta D(k_m) / \Delta R(k_m)$. Furthermore, since the link Lagrange multiplier $\Lambda_{(i,j)}$ is already available locally at each peer, then packets with a gradient $\Delta D / \Delta R$ greater than $\Lambda_{(i,j)}$ are forwarded. Otherwise, they are simply dropped from the local buffer and they are not forwarded in the overlay. Therefore, the problem for each peer is reduced to forwarding or dropping

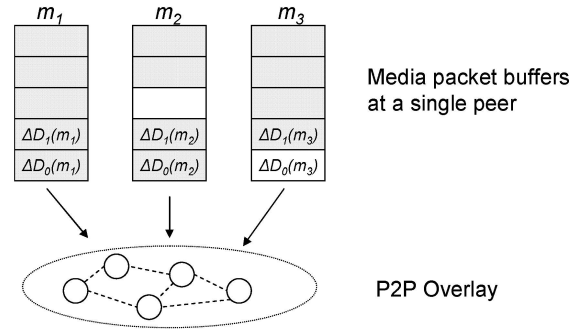


Fig. 2. Each peer maintains multiple buffers for the incoming media flows. Media packets are propagated in the overlay depending on the result of the rate allocation algorithm.

the packets depending on their relative importance that is expressed through their RD characteristics.

V. PERFORMANCE EVALUATION

In this section we present simulation results that compare the performance of different packet scheduling algorithms that are employed at the participating peers. We implemented the proposed algorithms as part of the ns-2 packet-level simulator. TCP with selective acknowledgments (SACK) is used for data transfer between the peers of the overlay. We tested peers with different upload capacities while we assumed that their download speed is 8 times higher in order to represent residential users with asymmetric last mile access links. Regarding the creation of the overlay network, the main principles we explained in Section II. More specifically, each peer requests from the central server a list of five peers from which it can receive content. This list is populated so as to distribute evenly the incoming requests. Therefore, no particular overlay maintenance algorithm is employed.

We compare the performance of the proposed RaDiO streaming algorithm, an earliest deadline first (EDF) scheduler, and a greedy algorithm that applies no particular scheduling technique (NO-SCHED). Instead of using actual video data, we created artificial and pre-formatted media content [13]. This form of artificial data units resemble a sequence of video frames that are temporally equidistant while they are also characterized by dependencies between them. By adopting this simulation approach the main advantage is that it allows us to focus precisely on the properties of the the underlying streaming algorithms which is the focus of this paper. Frames are composed of data units that are organized hierarchically depending on their importance for the receiver. All data units have the same size, set to 200 bytes. The increase in quality associated with a data unit is defined in units of quality. For example units that correspond to the same Intra-coded frame have an distortion value of 10. The quality by successful on-time reception of the media units is measured at each peer. Finally, the frame rate of the pre-formatted data is set to 30 frames per second while the initial pre-roll delay is 4 seconds.

A. Results

For our evaluation, we investigate the performance of the proposed streaming system in a static network that consists of 50 overlay nodes and two delivered flows. In Fig. 3(a), we present the average quality that is measured over all the clients in the overlay when one initial source (seed) of the data is employed in the experiment. We assume that a portion of these peers willingly forward/seed packets of a media flow. This represents more realistic conditions since the number of peers that altruistically contribute their resources is usually low. The important observation from the results in this figure is the ability of the RaDiO scheme to achieve improved quality by utilizing a reduced portion of the available channel capacity. The reason is of course that the RaDiO scheduler ensures that the data packets that are needed for successful decoding of subsequent packets are propagated throughout the overlay. Of course with sufficient additional initial pre-roll delay and increase in the available local memory resources at each peer, the delivered quality can be improved.

In Fig. 3(b), we observe similar results for the same number of overlay nodes but an increased number of peers that seed after they have received parts of the media description (50% of the population). As expected, the optimal quality is possible with a lower requirement on the upload capacity of a peer for all the scheduling schemes under test. Note that packets are propagated in the overlay even if they arrive late for their prescribed playback deadline since they may be used by subsequent packets for decoding.

VI. CONCLUSIONS

In this paper we presented a framework for RD-optimized media streaming in P2P overlay networks. Our scheme allows peers to forward through the overlay network media packets that have the highest impact on the quality of the reconstructed media sequence. This is accomplished by an optimization step that executes distributed rate allocation on a system-wide scale. Subsequently, peers enforce the calculated optimal rates by either forwarding or dropping packets. Our simulation results indicate significant performance improvement over techniques that do not consider the precise RD characteristics of media packet during the delivery in a P2P overlay network.

REFERENCES

- [1] D. Jurca, J. Chakareski, J.-P. Wagner, and P. Frossard, "Enabling adaptive video streaming in p2p systems," *IEEE Communications Magazine*, vol. 45, no. 6, pp. 108–114, June 2007.
- [2] W. Yiu, X. Jin, and S. Chan, "Challenges and approaches in large-scale peer-to-peer media streaming," *IEEE Multimedia Magazine*, vol. 14, no. 2, pp. 50–59, April-June 2007.
- [3] X. Zhang, J. Liu, B. Li, and T. P. Yum, "Coolstreaming/donet: A data-driven overlay network for live media streaming," in *IEEE INFOCOM*, March 2005.
- [4] N. Magharei, R. Rejaie, and Y. Guo, "Mesh or multiple-tree: A comparative study of live p2p streaming approaches," in *INFOCOM*, 2007, pp. 1424–1432.
- [5] M. Castro, P. Druschel, A. Kermarrec, A. Nandi, A. Rowstron, and A. Singh, "Splitstream: High-bandwidth multicast in a cooperative environment," in *SOSP*, October 2003.
- [6] N. Sarshar and X. Wuh, "Rate-distortion optimized network communication," in *INFOCOM*, May 2007.

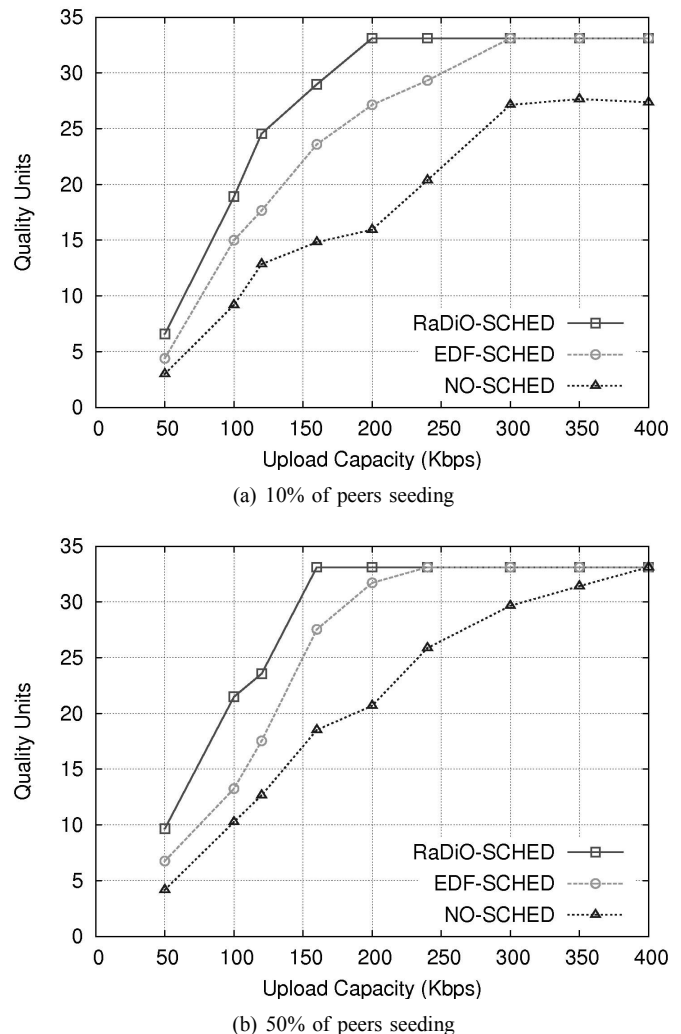


Fig. 3. Average number of quality units delivered on time for a peer population with 50 nodes.

- [7] P. A. Chou and Z. Miao, "Rate-distortion optimized streaming of packetized media," Tech. Rep., 2001.
- [8] R. Mahalingam, W. Tu, and E. Steinbach, "RD-optimized rate shaping for multiple scalable video streams," in *IEEE International Conference on Multimedia and Expo*, July 2007.
- [9] J. Chakareski and P. Chou, "Radio edge: Rate-distortion optimized proxy-driven streaming from the network edge," *IEEE/ACM Transactions on Networking*, vol. 14, no. 6, pp. 1302–1312, December 2006.
- [10] A. Argyriou, "Distributed resource allocation for network-supported FGS video streaming," in *Packet Video Workshop*, Lausanne, Switzerland, November 2007.
- [11] S. Low and D. Lapsley, "Optimization flow control-i: Basic algorithm and convergence," *IEEE Transactions on Networking*, vol. 7, no. 6, pp. 861–874, 1999.
- [12] D. P. Bertsekas and R. G. Gallager, *Data Networks*. Prentice-Hall, 1987.
- [13] C. D. Vleeschouwer, J. Chakareski, and P. Frossard, "The virtue of patience in low-complexity scheduling of packetized media with feedback," *IEEE Transactions on Multimedia*, vol. 9, no. 2, pp. 348–365, February 2007.