

Distributed Rate Allocation for Multi-Flow Video Delivery

Jacob Chakareski and Antonios Argyriou

ABSTRACT

We consider rate-distortion (RD) optimized multi-flow video delivery in unstructured overlay networks. We show that this problem can be studied as a distributed rate allocation. To solve the problem over the participating peers in the overlay, we apply classical decomposition techniques such that the network-wide utility of video distortion is minimized. Media packets are assumed to be piggy-backed with RD preambles that contain information regarding their impact on decoder video distortion and their size. This allows for converting the calculated optimal rate allocation at every node into simple forwarding or dropping actions. Furthermore, the proposed distributed media streaming framework employs a network inference algorithm for minimizing the flow of duplicate packets over the network and utilizing thus more efficiently the available resources. Our simulation results indicate that significant quality benefits can be achieved when the precise RD characteristics of a media presentation are taken into account.

1. INTRODUCTION

In recent years, the landscape of multimedia content distribution over the Internet has seen dramatic changes and considerable improvements. Initially, content delivery networks (CDN) and at a later stage peer-to-peer (P2P) networks, have emerged as the dominant solutions for point-to-multipoint video distribution in IP networks.¹ Both CDNs and P2P networks are based on the same fundamental concept, that of an overlay network. An overlay network is basically a collection of unicast sessions across cooperative network nodes with certain buffering and media processing capabilities.

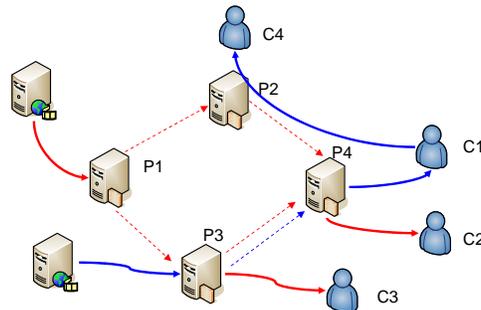


Figure 1. Media distribution through a hybrid P2P-CDN overlay network. Node P4 forwards the packets that maximize the quality at client C1 while avoiding receiving duplicate packets from nodes P2 and P3.

One of the first proposals that targeted efficient data delivery over multiple receivers in the absence of IP multicast is tree-structured application layer overlays.² However, for content distribution systems that employ either a dedicated application-layer overlay, or a P2P scheme that is characterized by dynamic node behavior, the tree structure has been shown to be problematic. Recently however, we have seen the advent of mesh-based overlay protocols that enable data dissemination between peers with no support from a regular overlay structure.³ In a typical gossip process, a node randomly selects a subset of target nodes to deliver recently received media packets, while it receives segments pushed from these nodes. In,⁴ the authors investigated the relative tradeoffs between three key parameters in live P2P streaming applications: server bandwidth cost, the maximum number of peers that can be supported, and the maximum number of streaming hops experienced by an active node. The problem of overlay bandwidth allocation has been addressed from the perspective of achieving service differentiation between different peers in.⁵

Nevertheless, two aspects of media distribution through the overlay model have been neglected. In most of the existing schemes the level of abstraction at which media flows are considered only deals with the delay, bandwidth

or other QoS requirements related to the nature of the media data. However, the precise impact of particular media packets on the quality of a decoded media stream, has not been taken into account. We believe that taking into account the RD characteristics of media packets is an important tool for optimizing the performance of large scale media streaming systems. However, for employing distortion-optimized streaming techniques, the particular challenges that are introduced when we consider thousands or even millions of nodes must be taken carefully into account. We address this need for scalability with the present paper by not adopting a full-fledged *RadIo* scheme,^{6,7} but instead we allow media packets to be self-contained with RD information so that they are easily manipulated in the overlay network. In summary, this paper presents a novel algorithm for distributed rate allocation algorithm for multi-flow video distribution in overlay networks that works in conjunction with a distributed algorithm for packet duplicate suppression that we also design.

The rest of this paper is organized as follows. In Section 2, we present our system model and motivate the design of our system. We define formally the optimization problem in Section 3. The distributed and centralized algorithms are presented in Section 4. Next, we present simulation results in Section 5. Finally, Section 6 presents our conclusions and provides possible directions for future work.

2. SYSTEM MODEL

We assume a mesh-based P2P delivery model where participating peers form a randomly connected and directed mesh. Transport layer connections between peers are 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. For the initial part of our analysis we assume that an active node can obtain a list of the currently participating active nodes in the overlay from a central server (a.k.a. Tracker). This bootstrapping server maintains a list of all participating nodes and provides a random subset of participants to each new peer. Each of the participating active nodes delivers a subset of the *media sessions*. In a P2P network, usually an active node is receiving data for a single media flow and so it can contribute part of its uplink capacity only for this flow. But in the general case of a CDN this is not the case since the active node might need to allocate the available resources between several flows. Regardless of the precise functionality of each active node, the overlay topology can be modeled as a directed acyclic graph. Let us denote this DAG as $G = \{N, A\}$, where N is the set of all overlay nodes while A represents the connections between them. Furthermore, M denotes the number of media flows currently being delivered by the overlay while A_m is the set of links between peers that are used for delivering media flow $m = 1, \dots, M$. We denote the index of a media packet from media stream m as k_m . Similarly to related works, e.g.,⁶ 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 denoted as $\Delta D(k_m)$. In practice $\Delta D(k_m)$ is the total increase in mean square error (MSE) distortion affecting the video stream if the packet is not delivered to its destination by its prescribed deadline.⁶

3. PROBLEM FORMULATION

For an active node the goal is simple: Calculate the optimal bandwidth allocation to incoming media flows so that the quality at the receiving clients is maximized. We are interested in running this optimization in collaboration with other nodes such that a network-wide utility is maximized. In the following, we show that this is a distributed rate allocation problem that can be formulated as a constrained optimization. However, before proceeding with the formal definition of the aforementioned problem, we introduce further necessary notation first. Let μ_i denote the group of incoming flows at node i . Also 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 $r_m^{(i,j)}$. Then, the overall distortion of all flows being delivered over the link $i \rightarrow j$ 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 a weight factor that node i can assign to media flow m when performing rate allocation over the outgoing link (i, j) . This weight factor can express the importance that the original sender of flow m assigns to this flow in terms of its delivery requirements.

Using (1), we can express the average media distortion over all links in the overlay as:

$$D_N = \sum_{(i,j) \in A} D^{(i,j)}. \quad (2)$$

We are interested in minimizing the network distortion D_N subject to the available bandwidth $R^{(i,j)}$ at each link $(i,j) \in A$. Recall that $r_m^{(i,j)}$ denotes the rate assigned to flow $m \in \mu_i$ on link (i,j) , as introduced above. Then, $r^{(i,j)} = \sum_{m \in \mu_i} r^{(m)}(i,j)$ denotes the overall rate of the flows sent over this link. Hence, the optimization problem under consideration can formally be written as:

$$\begin{aligned} & \min D_N, \\ \text{s.t. } & r^{(i,j)} \leq R^{(i,j)}, \quad \forall (i,j) \in A. \end{aligned} \quad (3)$$

4. DISTRIBUTED SOLUTION

4.1 Global Rate Allocation

We apply Lagrange duality to the constraint in (3), and produce the partial Lagrangian as follows:

$$L = D_N + \sum_{(i,j) \in A} \Lambda_{(i,j)} r^{(i,j)}, \quad (4)$$

where $\Lambda_{(i,j)} > 0$ is the Lagrange multiplier for link (i,j) at active node i . The dual function is then defined as

$$g(\Lambda) = \max_{\Lambda > 0} L(\Lambda), \quad (5)$$

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

$$\begin{aligned} & \min g(\Lambda) \\ & \text{subject to } \Lambda \geq 0 \end{aligned} \quad (6)$$

Note that the Lagrange multiplier sets 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 (3). Furthermore, this Lagrange multiplier decomposes the primal problem into individual streaming rate allocation problems at every active node that can be optimized individually. In particular, each active node i must solve for the optimal rate allocation:

$$r^{*(i,j)} = \arg \min \left\{ D^{(i,j)} + \Lambda_{(i,j)} r^{(i,j)} \right\}, \quad (7)$$

The convergence of the dual algorithm to the optimal solution has been proven for distributed flow control problems with a convex utility function.⁸ For calculating the Lagrange multiplier $\Lambda_{(i,j)}$ at every node i , we employ a gradient method⁹ that updates $\Lambda_{(i,j)}$ as

$$\Lambda_{(i,j)} = \max \left\{ 0, \Lambda_{(i,j)} - \beta \left(R^{(i,j)} - r^{*(i,j)} \right) \right\} \quad (8)$$

every time after rate allocation is performed, where β is a small constant.

4.2 Multi-Flow Rate Allocation at a Single Active Node

At the next level of the optimization, the optimal rate allocation is calculated for the flows that are sharing a common outbound link at an active node. This is done based on the importance of the forwarded media flows and reflected through the gradient factor $\lambda(m) = dD_m(R)/dR$. 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 $\gamma_m dD_m(R)/dR \approx \Lambda_{(i,j)}$. The rate value associated with this point is the optimal value that should be assigned by node i to the incoming flow m on the outgoing link j . Let this value be denoted as $r_m^{*(i,j)}$.

Since a media flow is already assigned an optimal outgoing rate $r_m^{*(i,j)}$, an active node can easily allocate this rate between the individual packets that are locally stored and belong to that flow. Recall from Section 2 that information on the rate-distortion characterization is piggy-backed with each individual packet. Each packet is labeled with ΔD and ΔR . This information provides the per-packet gradient $\Delta D/\Delta R$. Furthermore, since the link Lagrange multiplier $\Lambda_{(i,j)}$ is already available locally at each peer, packets with a gradient $\Delta D/\Delta R$ greater than $\Lambda_{(i,j)}$ are forwarded by the node. Otherwise, they are simply dropped from the local buffer and are not forwarded further.

5. PERFORMANCE EVALUATION

In this section, we present simulation results that compare the performance of different packet scheduling algorithms employed by the active nodes. We implemented the proposed algorithms as part of the ns-2 packet-level simulator. TCP with selective acknowledgments (SACK) is employed for data transfer between the peers of the overlay. We tested active nodes 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. We assume nodes join the overlay and forward content while they are online. The nodes depart the network at the end of their playback. Each active node requests from the central server a list of 5 active nodes from which it can receive content. This list is populated so as to distribute evenly the incoming requests. However, no particular overlay maintenance algorithm is employed since upon the departure of an a node the central server is contacted again for querying a replacement neighbour.

The systems under comparison are the proposed *RaDiO* streaming algorithm, an earliest deadline first (EDF) scheduler, and a greedy algorithm that applies no particular media-aware scheduling technique (*NO-SCHED*). Essentially what we are comparing with these experiments are different chunk/packet selection strategies. Instead of using actual video data, we created artificial and pre-formatted media content.¹⁰ 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 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 integer units of quality. For example, units that correspond to the same Intra-coded frame have a 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 fps while the initial pre-roll delay is 4 seconds.

5.1 Varying the Upload Capacity

Here, we investigate the performance of the three systems in a static network comprising 50 overlay nodes and one flow to be delivered. 10% of the nodes act as seeding servers to the rest of the peer population. In other words, these nodes have initially the content at their disposal and serve it throughout the duration of the simulation to their neighbours in the overlay topology. In the experiment, we measure media quality at every node as a function of the upload bandwidth of the peers. In Figure 2, we show the quality averaged over all nodes in the overlay. The important observation from these results is the ability of the *RaDiO* scheme to achieve improved quality by efficiently utilizing a reduced portion of the available channel capacity. This is due to the fact that the *RaDiO* scheduler ensures that the data packets that are needed for successful decoding of subsequent packets are propagated throughout the overlay first.

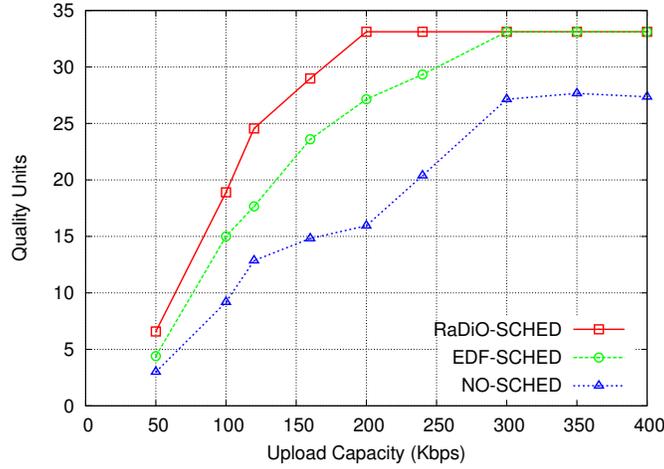


Figure 2. Average number of quality units delivered on time (10% seeding nodes).

In Figure 3, we observe similar results for the same number of overlay nodes but an increased number of seeding nodes (50%). For simplifying the simulation process, in order to increase the server bandwidth in the network, in these experiments we opted to increase the number of servers while keeping their upload capacity same with the rest of the nodes in the network, as in Figure 2. As expected, achieving a certain quality is possible now with a lower requirement on the upload capacity of an active node for all the scheduling schemes under examination. 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.

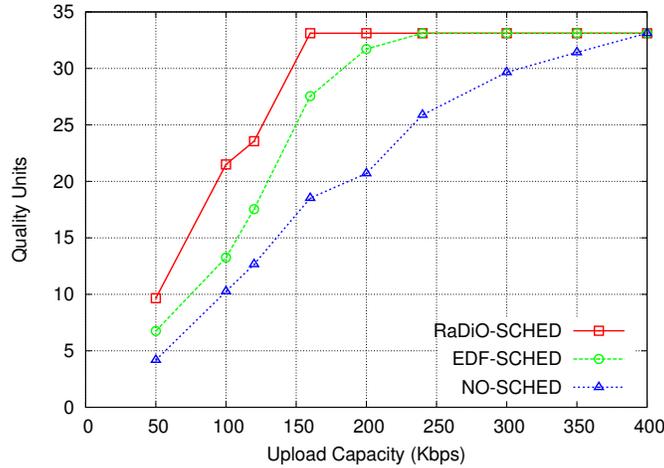


Figure 3. Average number of quality units delivered on time (50% seeding nodes).

5.2 Server Percentage in a CDN-P2P Topology

In Figure 4, we examine the influence of the number of seeding servers in the network on the resulting quality at every peer. As in the previous two sets of experiments, the servers have the same upload capacity as the rest of the nodes in the network and their number is expressed as a percentage of the total peer population. The rationale behind this experiment corresponds to a hybrid CDN-P2P network where we want to identify the critical number of seeds/servers that initially need to have the content in order to provide a satisfactory quality at each peer. In the experiment, the peer population is 300 nodes and the upload capacity at each node is

considerably higher now at 1000 kbps. As we can see from Figure 4, for the same average quality a much smaller number of initial seeds or servers is needed when the *RaDiO-SCHEd* scheme is deployed. For example, for average quality of 30, we achieve a 50% reduction in number of servers (server capacity) required in the network relative to *EDF* scheduling.

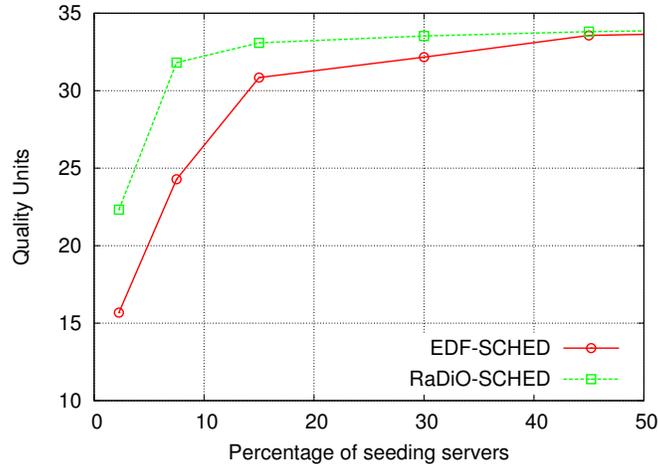


Figure 4. Average number of quality units delivered on time (300 nodes with 1000 kbps upload capacity each).

6. CONCLUSIONS

We presented a framework for RD-optimized media streaming in P2P overlay networks. Our scheme allows peers to forward through the 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 packets or not. Our simulation results indicate significant performance improvement over techniques that do not consider the RD characteristics of media packets during their delivery in P2P overlay networks. In the full paper, we will additionally investigate the performance of the proposed framework over other performance metrics and we will provide further detail on the optimization performed at every active node.

REFERENCES

- [1] Jurca, D., Chakareski, J., Wagner, J.-P., and Frossard, P., “Enabling adaptive video streaming in p2p systems,” *IEEE Communications Magazine* **45**, 108–114 (June 2007).
- [2] Chu, Y., Rao, S., and Zhang, H., “A case for end system multicast,” in [*ACM SIGMETRICS*], (2000).
- [3] Zhang, X., Liu, J., Li, B., and Yum, T. P., “Coolstreaming/donet: A data-driven overlay network for live media streaming,” in [*IEEE INFOCOM*], (March 2005).
- [4] Small, T., Liang, B., and Li, B., “Scaling laws and tradeoffs in peer-to-peer live multimedia streaming,” in [*ACM Multimedia*], (October 2006).
- [5] Wu, C. and Li, B., “Diverse: Application-layer service differentiation in peer-to-peer communications,” *IEEE Journal on Selected Areas in Communication* **25**, 222 – 234 (January 2007).
- [6] Chou, P. A. and Miao, Z., “Rate-distortion optimized streaming of packetized media,” *Microsoft Research Technical Report MSR-TR-2001-35* (2001).
- [7] Chakareski, J. and Chou, P., “Radio edge: Rate-distortion optimized proxy-driven streaming from the network edge,” *IEEE/ACM Transactions on Networking* **14**, 1302–1312 (December 2006).
- [8] Low, S. and Lapsley, D., “Optimization flow control-i: Basic algorithm and convergence,” *IEEE Transactions on Networking* **7**(6), 861–874 (1999).
- [9] Bertsekas, D. P. and Gallager, R. G., [*Data Networks*], Prentice-Hall (1987).

- [10] Vleeschouwer, C. D., Chakareski, J., and Frossard, P., “The virtue of patience in low-complexity scheduling of packetized media with feedback,” *IEEE Transactions on Multimedia* **9**, 348–365 (February 2007).