Computer Networks 91 (2015) 483–494

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Optimal bandwidth allocation for hybrid Video-on-Demand streaming with a distributed max flow algorithm

Chen Tian[a], Jingdong Sun[a], Weimin Wu[a,*], Yan Luo[b]

[a] School of Electronic Information and Communications, Huazhong University of Science and Technology, 1037 Luoyu Road, Wuhan, China
[b] Department of Electrical and Computer Engineering, University of Massachusetts Lowell, Lowell, MA, USA

ARTICLE INFO

Article history:
Received 29 October 2014
Revised 2 August 2015
Accepted 27 August 2015
Available online 11 September 2015

Keywords:
Video-on-Demand
Streaming
Bandwidth allocation
Max flow
Distributed algorithm

ABSTRACT

Video-on-Demand (VoD) streaming counts for the largest share of Internet traffic, and commonly relies on P2P–CDN hybrid systems. In such hybrid systems, a peer's upload bandwidth utilization is critical to P2P mode, thus the bandwidth allocation algorithm is important to the performance of VoD streaming. The current research either makes impractical assumptions, or is inefficient in real scenarios. Meanwhile some industrial practice such as additive-increase/multiplicative-decrease (AIMD) heuristic, although effective in practice, lacks theoretical foundation. This paper develops an optimal bandwidth allocation algorithm for hybrid VoD streaming. Specifically, we propose a novel *Demand Driven Max-Flow Formulation*, which treats each peer's bandwidth demand as the flow commodity. The proposed distributed *Free-for-All Push–Lift* algorithm can solve the formulation in each peer, and is free of any lock, shared memory and atomic operation. Following the theoretical analysis, we implement the algorithms in real-world VoD streaming systems. Through extensive evaluations we demonstrate that our approach can provide high-quality bandwidth allocation for hybrid VoD systems.

1. Introduction

Motivation Online streaming has been dominating the Internet traffic [1] and Video-on-Demand (VoD) is the primary method of content consumption [2,3]. Among VoD services, some services such as NetFlix and Hulu rely purely on Content Distribution Networks (CDNs) while others exploit Peer-to-Peer (P2P) transport to reduce service cost [4–6]. A growing trend is however leveraging P2P–CDN hybrid systems [5–10] for VoD services, where the contents requested by a user are delivered by either her peers or the CDN . Most modern P2P streaming networks are operated in such a hybrid mode. One representative example is the RTMFP-based P2P support of Adobe Flash platform [11]. In such hybrid platforms, it is beneficial to maximize the content distribution through peers rather than CDN due to CDN hosting costs.

The upload bandwidth utilization of peers is critical for VoD in P2P mode. Each peer has two roles: uploader and downloader. The more upload bandwidth a peer could exploit from its neighboring peers, the more data it could download from them. For hybrid systems, more downloaded data from peers translates to fewer video piece missing events in the P2P mode [12,13], hence reducing data serviced from CDN servers.

However, the existing approaches are ineffective or inadequate in optimizing the upload bandwidth utilization for peers. Many of the prior research either make impractical assumptions, or are inefficient in real-world scenarios. Some works such as [14,15] assume a complete connected graph, which is not realistic in real P2P systems. Zhao's work [16] showed that with a random peer selection and uniform bandwidth allocation, a system can asymptotically achieve a close to the optimal streaming rate if each peer

* Corresponding author. Tel.: +862787543236.
E-mail addresses: tianchen@hust.edu.cn (C. Tian), jdsun@hust.edu.cn (J. Sun), wuwm@hust.edu.cn, 565714801@qq.com (W. Wu), yan_luo@uml.edu (Y. Luo).

http://dx.doi.org/10.1016/j.comnet.2015.08.049
1389-1286/

maintains $O(\log N)$ downstream neighbors. For every practical system, there is an upper bound of the number of neighbors; there are also many scenarios where the uniform rate allocation certainly cannot work efficiently (see Section 3 for details). To the best of our knowledge, the common industry practice is based on the additive-increase/multiplicative-decrease (AIMD) heuristic. Motivated by TCP congestion control [17], AIMD is a simple allocation protocol in the P2P overlay layer [2,18]. Although proved effective in practice, AIMD lacks theoretical foundations thus is not extensible.

We are motivated to model the bandwidth allocation problem in hybrid VoD platforms and provide efficient solutions. Central to the upload bandwidth utilization is the bandwidth allocation algorithm, which decides in each peer how its bandwidth should be allocated to its neighboring peers. We aim to design such a bandwidth allocation algorithm that satisfies two important criteria. First of all, a good algorithm should not leave (much) unused bandwidth. Due to the pressure of market competition, the video streaming rates become higher and higher; usually in a video channel, peers' aggregated upload bandwidth is either almost equal to or less than the aggregated steaming bandwidth demand. It is beneficial to maximize the overall bandwidth utilization. Second, a good algorithm should allocate to each peer *just enough* bandwidth that can support the streaming rate. For a single peer, if the aggregated allocated bandwidth is less than the rate, it would download a significant portion of data from CDN, which needs to be prevented. The allocated bandwidth should not be (much) larger than the streaming rate either: each peer is consuming the video at the same rate; a peer exploiting unnecessary bandwidth is just accumulating its own buffered data. If the peer quits the channel early or skips the segment, which are very common [5], all downloaded data are wasted. In summary, unnecessary bandwidth in one peer would eventually reduce the download rates of others, which in turn deteriorates the performance of the whole system. In designing such an algorithm, we would like to ensure that the approach is guided by theoretical analysis, and implemented under practical constraints of real VoD systems.

Challenge In practice it is challenging to design an optimal bandwidth allocation algorithm as explained below.

The first challenge is *how to formulate the optimal objective.* To maximize peers' contribution, a max-flow formulation is the natural choice. The real question is: directly modeling allocated bandwidth between a pair of peers as flow commodity is intuitive, while not suited for the receiver-driven nature of streaming systems (see Section 4 for details).

The second challenge is *how to solve the optimal formulation.* Bandwidth allocation has to be solved in a distributed manner. Centralized scheduling is not applicable in P2P systems: the topology of peer connections is inherently dynamic, while a centralized control has to re-calculate the max-flow problem each time a peer joins or leaves the channel and then has to distribute the results back to every peer.

The third challenge is *how to convert the theoretical design to a practical implementation.* The real P2P scenarios have many requirements: *e.g.*, after agreement of bandwidth allocation between an uploader and a downloader, how to realize it in an uploader? What if the promised bandwidth is temporarily larger or lower than the actual bandwidth? How to tolerate topology dynamics?

Contribution We make the following contributions in this paper to address the challenges:

- We propose a *Demand Driven Max-Flow Formulation*, which treats each peer's bandwidth demand as flow commodity. This formulation is friendly to the receiver-driven nature of streaming systems (Section 4).
- We propose a distributed *Free-for-All Push–Lift* algorithm to solve the optimization formulation and overcomes real distributed system constraints, such as the lack of global locks, shared memory and atomic operations (Section 5).
- We transform the theoretical design into a practical implementation by meeting real system requirements, such as piece scheduling and topology dynamics (Section 6).
- We conduct extensive evaluations and demonstrate that the proposed approach can provide high quality bandwidth allocation for hybrid VoD systems. In every scenario, its performance is better than AIMD (Section 7).

The rest of paper is organized as follows. Section 2 discusses the related work. Section 3 gives the background of hybrid VoD systems and the bandwidth allocation problem. Section 8 concludes the paper.

2. Related work

There are two previous studies closely related to our research. Although Yu and Chen model bandwidth allocation in P2P systems as a max-flow problem, they use a *Supply-Driven* formulation [19]. As mentioned in Section 4, *Supply-Driven* is not well suited for distributed receiver-driven nature of P2P systems. Also, their work solves the formulation using a centralized multi-stage max-flow algorithm. We have pointed out that centralized scheduling is not applicable in P2P systems, given that the topology is highly dynamic. Therefore, Yu's work of applying centralized computation on a static graph is impractical for hybrid streaming networks. He et al. present a systematic study on the throughput maximization problem in P2P–VoD applications [20]. A fully distributed algorithm is proposed for both buffer–forwarding and hybrid–forwarding architectures. The basis of their algorithm is linear programming with Lagrangian duality. However, any changes in peer's upload or download bandwidth requires a new programming execution on the whole graph. It is also thus impractical for highly dynamic P2P systems.

There are research studies on bandwidth allocation for multi-tree based overlay networks [21,22], which are related to P2P based VoD streaming. Our paper focuses on mesh-based hybrid VoD networks without content bottlenecks. To our knowledge, this is the first paper that presents a proved distributed algorithm, and implements a practical design for such systems. Bradai et al. focus on bandwidth allocation and the incentive mechanisms for layered video streaming [23]; our paper instead focuses on the widely supported non-layered VoD systems.

Some research papers also involves bandwidth allocation but in the scope of live streaming. Wu et al. propose bandwidth allocation algorithms to properly provision bandwidth among multiple live streaming channels [24]. ShadowStream project directly uses the industry AIMD approach [12,13]. One of our previous work also deals with topology dynamics such