PEER-TO-PEER VIDEO ON DEMAND: CHALLENGES AND SOLUTIONS

Vijay S. Pai[*], *Yung Ryn Choe*[†], *Jagadeesh Dyaberi, Derek Schuff, Karthik Kannan*

Purdue University
West Lafayette, IN 47907

ABSTRACT

The challenges and solutions required for peer-to-peer video-on-demand (P2P VoD) provided by a fixed provider such as a cable company are fundamentally different from those seen in traditional P2P networks or client-server VoD solutions. Unlike traditional P2P networks, the end nodes (set top boxes with DVR capabilities) are largely under control of the system provider. Consequently, issues like churn and free-loading are less substantial. Unlike client-server solutions, there is always a readily-available resource of peer nodes able to contribute *even if they are not using the VoD service!*

This paper explores requirements for efficient preloading of VoD movie data onto numerous customer set-top boxes. This research is currently exploring mathematical programming algorithms that minimize uplink traffic, given a popularity model for various pieces of content and information about storage and bandwidth capacity constraints at the customer nodes. Given the complex non-linear nature of P2P interactions, these mathematical programs require non-linear optimization approaches or heuristic solutions. However, even heuristic solutions would likely provide substantial advantages over simple dynamic allocation.

Index Terms— Video-on-demand, mathematical programming

1. INTRODUCTION

As systems such as voice, video, and data communication merge onto a single IP-network delivery platform, users are likely to benefit with a greater number of choices and more content control at their disposal [1]. Given the numerous advantages of doing so, cable companies have made significant investments in systems that seek to deliver high-bandwidth digital content (e.g., Video-On-Demand service for full length, high-definition movies) to consumers. However in the current unicast data delivery model, all the movie data is streamed from the regional or national-level video hubs to the requesting viewers. This unicast data delivery model puts a considerable amount of strain on the regional/national servers and on the uplink connection between the local community switch and the local video hubs. The uplinks between the communities and the server hubs thus become the bottleneck in determining the total number of concurrent viewers, and increasing this number will only lead to more usable systems for the users and greater potential revenue for the network providers.

2. PROPOSED SOLUTION AND CHALLENGES

The viewers in this environment are users of set-top boxes provided by the networks. Although set-top boxes can store viewer-selected content through resources such as DVR, they are essentially largely under control of the system provider. Thus, the system provider may use their resources as a way of extending the capabilities of the network. A feasible solution to reduce uplink load and extend the network is to use P2P communication among set-top boxes in a community. This approach could exploit any unused bandwidth in the community network and any unused capacity in the user set-top boxes.

In our previous work we modified the BitTorrent protocol to enable peers (set-top boxes) to stream data to each other [2]. Our results showed that with a modified piece-picker policy, data demand on the server is reduced by upto 70-90% when compared to simple unicast delivery model, thus validating the use of P2P communication for streaming. Our system was evaluated with a single movie with no constraints placed on the piece-popularity, as this is most similar to the existing BitTorrent file swarm model [3]. However, using individual swarms for videos would mismanage and/or underutilize the potential resources available in a network of set-top boxes.

We now plan to further our research into P2P communication for video content delivery by extending our system to consider the streaming of all movies available from the VoD provider. We have identified several research problems that arise by considering an extended system.

First, porting a traditional P2P model into a environment that is largely centrally-controlled would underutilize the resources. If end nodes are entirely responsible for making decisions about which content to store based on their own on-demand viewing, the system will not be increasing its capabil-

[*]This work is supported in part by the National Science Foundation under Grant Nos. CCF-0532448 and CCF-0621457.

[†]Currently at Sandia Livermore National Labs

ities whenever end nodes are idle. This is a substantial issue in situations like home VoD where we would expect idleness for most of the day followed by a brief period of high activity in the evening and early night. Instead, the system should try to preload desired data onto the set-top boxes and thus make all set-top boxes active participants in the P2P network even if they are not consuming P2P resources.

Second, the issue of set-top box capacity becomes a concern when we consider all the content available from the provider (which can reach thousands of movies). The total capacity among all the set-top boxes in a community is far less than the capacity needed to store all the movies, so only a subset of the total content can be stored in the community. Further, if the movies are encoded in high-definition, capacity becomes an even bigger constraint.

Third, bandwidth is closely tied with capacity as both must be considered jointly to maximize overall streaming efficiency. This is a concern as all networks currently exhibit substantial asymmetry between the uplink and downlink bitrates at any given end node (typically about an order of magnitude). If each end node has an uplink of 1 Mbps, it would require six sources simultaneously contributing data to satisfy a single request for a HD-quality streaming movie encoded at 6 Mbps; each of these sources would be fully saturating its uplink. This issue now relates back to capacity, because offloading a substantial amount of streaming traffic onto the P2P network requires that enough copies of the data exist in the network to account for the asymmetry in uplink and downlink rates. This requires replication, which reduces the effective aggregate capacity in the network and makes the subset of content that can be stored in the network an even smaller fraction of the total.

Because the P2P network can only store a subset of the data content available, any preloading solution must intelligent choose which data to load and replicate. An effective decision would require an understanding of the popularity of different pieces of content, as the greatest benefits would arise from replicating the most popular content. At the same time, the high demand for these most popular videos would require extensive replication to be able to support as many simultaneous requests as are likely to arrive.

We are currently exploring different intelligent capacity allocation policies using mathematical models that would seek to maximize the streaming availability and efficiency within the community network. Putting together capacity and bandwidth constraints, we are developing mathematical programming models that build up a "presence matrix" of pieces of content by carefully considering capacity constraints (making sure that no given node is assigned too much data) and bandwidth constraints (making sure that no node's uplink is overly strained by heavily replicating the most popular pieces of content and distributing the popular content across different sets of nodes so that no node is particularly more demanded than others) while optimizing for an objective function. The current objective function is to minimize the aggregate bandwidth demand at the servers.

3. SUMMARY AND FUTURE WORK

Although P2P VoD offers substantial performance challenges, intelligent preloading algorithms should enable effective utilization of system resources to allow scalability beyond the saturation point of the current client-server model. Consequently, this should be an important area for short-term and long-term contributions.

Future work must consider the robustness and security of the proposed P2P delivery model. One of the biggest issues in traditional P2P research is churn: Nodes come and go randomly which changes the network dynamics. Capacity and bandwidth can change instantaneously and the network should be able to adapt to the change. Although less than a traditional P2P system, a set-top box network can still experience churn caused by power failures or users turning off the power on their node; additionally, a set-top box is in a state of "virtual churn" if its uplink is being used entirely for other purposes at a given time. We would like to explore robustness and resilience against churn in the system. Additionally, using set-top boxes to make decisions and serve data introduces an increased possibility of malicious users looking to manipulate the system to their own advantage. It is thus also important to protect content and detect and prevent harmful activity.

Although cable companies may deploy a setup like this for VoD, they can use it to gain additional benefits in other domains as data traffic is also carried on the same network as videos. Thus, the infrastructure that results from these efforts may benefit other applications. For example, this peer-to-peer setup may also act as a cache for Web content or blocks of other commonly- accessed files, allowing a general reduction in uplink traffic from non-VoD sources as well. Such efforts may be important when dealing with high-bandwidth events such as software updates that affect all end-users at once.

4. REFERENCES

[1] Yih-Farn Chen et al., "When is P2P Technology Beneficial for IPTV Services," in *Proceedings of the 17th International Workshop on Network and Operating System Support for Digital Audio and Video*, May 2007.

[2] Yung Ryn Choe, Derek L. Schuff, Jagadeesh M. Dyaberi, and Vijay S. Pai, "Improving vod server efficiency with bittorrent," in *MULTIMEDIA '07: Proceedings of the 15th international conference on Multimedia*, New York, NY, USA, 2007, pp. 117–126, ACM.

[3] Clive Thompson, "The BitTorrent Effect," *Wired Magazine*, January 2005.