Signal Processing: *Image Communication* 27 (2012) 343–359

Contents lists available at SciVerse ScienceDirect

Signal Processing: *Image Communication*

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

Watching user generated videos with prefetching

S. Khemmarat [a,*], R. Zhou [b], D.K. Krishnappa [a], L. Gao [a], M. Zink [a]

[a] *Department of Electrical and Computer Engineering, University of Massachusetts, Amherst, MA 01003, USA*
[b] *College of Computer Science and Technology, Harbin Engineering University, Harbin, China*

ARTICLE INFO

Available online 25 November 2011

Keywords:
Prefetching
User generated videos
Performance

ABSTRACT

Even though user generated video sharing sites are tremendously popular, the experience of the user watching videos is often unsatisfactory. Delays due to buffering before and during a video playback at a client are quite common. In this paper, we present a prefetching approach for user-generated video sharing sites like YouTube. We motivate the need for prefetching by performing a PlanetLab-based measurement demonstrating that video playback on YouTube is often unsatisfactory and introduce a series of prefetching schemes: (1) the conventional caching scheme, which caches all the videos that users have watched, (2) the search result-based prefetching scheme, which prefetches videos that are in the search results of users' search queries, and (3) the recommendation-aware prefetching scheme, which prefetches videos that are in the recommendation lists of the videos that users watch. We evaluate and compare the proposed schemes using user browsing pattern data collected from network measurement. We find that the recommendation-aware prefetching approach can achieve an overall hit ratio of up to 81%, while the hit ratio achieved by the caching scheme can only reach 40%. Thus, the recommendation-aware prefetching approach demonstrates strong potential for improving the playback quality at the client. In addition, we explore the trade-offs and feasibility of implementing recommendation-aware prefetching.

1. Introduction

The advent of user-generated video sharing sites such as YouTube, Dailymotion, Metacafe, Tudou, and Daum has provided tremendous opportunities for Internet users to share their personal experiences as well as to conduct business. The enormous amount of video content uploaded on video sharing sites has made these sites information sources, which Internet users often turn to for information, entertainment, and even education. For example, YouTube has hundreds of millions of viewers and delivers billions of videos each month. Unlike the traditional video-on-demand (VoD) systems that typically deliver professionally produced content, video sharing sites typically contain short video clips produced for a particular purpose [1]. The short duration of video clips combined with the huge collection of videos makes it possible for users to browse around for content of interest.

Despite the tremendous popularity of user generated video sharing sites, user experience with watching videos from these sites can vary significantly [2]. In this paper, we perform a PlanetLab-based measurement to investigate user experience with watching videos on YouTube by analyzing the number of pauses and pause duration during video playback. Our results demonstrate that it is common for a user to experience one or more pauses when watching a video on YouTube. These interruptions during video playback can be quite annoying and can potentially discourage users from watching more videos

* Corresponding author.
E-mail addresses: khemmarat@ecs.umass.edu (S. Khemmarat), renjie_zhou@hrbeu.edu.cn (R. Zhou), krishnappa@ecs.umass.edu (D.K. Krishnappa), lgao@ecs.umass.edu (L. Gao), zink@ecs.umass.edu (M. Zink).

0923-5965/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.image.2011.10.008

or simply turn users off at the very beginning of a video browsing session. Even a small number of pauses can have a very negative impact since the majority of videos on video sharing sites are usually relatively short (on the order of a few minutes) [3,4]. Clearly, an increase in network bandwidth and scalable solutions on video sharing sites can solve some of these problems. However, the desire for and the increasing availability of high quality videos (such as high quality or high definition videos) might further exacerbate the experience of browsing video sharing sites.

In this paper, we propose to prefetch video content in order to reduce or eliminate the potential of pauses during video playback and decrease playback delay. We introduce a series of prefetching schemes: conventional caching scheme, search result-based prefetching scheme, and recommendation aware-prefetching scheme. Our proposed prefetching scheme conserves bandwidth by prefetching only a *prefix* of a video since a video clip can playback smoothly if a sufficiently large prefix of the video is prefetched [5]. Furthermore, the prefetching scheme can take advantage of many "idle" periods of a video browsing session by prefetching when the current playback does not fully utilize the available bandwidth or when users read comments between watching videos.

We evaluate our proposed prefetching schemes with user browsing pattern data collected from a university network. We focus on user browsing patterns on YouTube since YouTube is the most popular video sharing website in North America. Our measurement results show that our recommendation-aware prefetching approach can achieve an overall hit ratio of up to 81%, while the hit ratio achieved by the caching scheme and the search result-based prefetching scheme can only reach 40% and 38%, respectively. Therefore, our study demonstrates strong potential for improving playback quality at the client using recommendation-aware prefetching. Although building an effective recommendation system itself is a challenge [6], it can potentially provide sufficient clues for predicting videos that are most likely to be watched next by users. We also explore the trade-offs and feasibility of implementing the recommendation-aware prefetching.

Although our evaluation is presented in the context of a proxy cache architecture, the proposed prefetching scheme can potentially be applied to a peer-to-peer architecture or to servers in content delivery networks (CDNs). Despite the fact that prefetching has been proposed in the context of Web and multimedia delivery, demonstrating its effectiveness has been challenging without user browsing traces. To the best of our knowledge, our work is the first to systematically measure and compare the effectiveness of various prefetching schemes based on *actual user browsing activities* and demonstrate the advantage of exploiting the recommendation system for video delivery.

The rest of the paper is organized as follows. In Section 2, we investigate the user experience on YouTube regarding the pauses users would experience in a video playout. Section 3 describes the prefetching schemes and the algorithms to select videos to prefetch. In Section 4, we describe our datasets and measurement of the sources of video views. The evaluation of the proposed prefetching schemes is presented in Section 5, and in Section 6 we discuss the trade-offs and feasibility of prefetching. Related work is presented in Section 7. Finally, Section 8 concludes the paper.

2. Investigating user experiences with watching YouTube videos

Previous work has shown that service delay on YouTube is longer than on other video sharing websites [2]. To further demonstrate the need for prefetching, we perform an experiment to evaluate user experience in watching YouTube videos. In particular, we measure how likely it is for a user to experience pauses during video playback and how long the pauses are. We first describe our data collection methodology and how we emulate the playback. Then, we present our results on estimating the possibility and duration of pauses experienced by a viewer.

2.1. Data collection

To conduct a large scale and world wide study on the user experience on YouTube, instead of requiring real users to watch videos and collect traces as we have done in [7], we investigate the user experience by analyzing *video download traces*. A video download trace is a trace of incoming and outgoing network traffic captured while a user is watching a video on YouTube. We utilized PlanetLab network in order to obtain video download traces from several locations in the world. For each PlanetLab node, we select a sample set of videos and automate the trace capturing of the videos' download. The video download traces are then analyzed for pause information by our tool. The combination of the use of PlanetLab network and the automated trace capturing and pause analysis allows us to investigate world wide user experience without requiring actual users to be present to watch videos. In the following subsections, we describe in detail our process of obtaining the video download traces.

2.1.1. PlanetLab nodes

PlanetLab [8] is a global research network that supports the development of network services like distributed storage, network protocols, and peer-to-peer systems. There are more than 1000 PlanetLab nodes at 500 distinct sites distributed all over the world.

For our measurement, we group PlanetLab nodes into three different regions based on their locations: USA (US), Europe (EU), and Asia (AS). The grouping allows us to analyze whether video playback quality differs based on the physical location a video is transmitted to. For our experiment we used 70 US nodes, 47 EU nodes, and 24 AS nodes. The PlanetLab nodes we used are distributed in several physical locations, including various universities and organizations, which means that our dataset contains sample points from diverse locations and networks in each region.