Watching User Generated Videos with Prefetching

Samamon Khemmarat, Renjie Zhou, Lixin Gao, Michael Zink

Electrical and Computer Engineering Department UMass Amherst

MMSys 2011, Feb. 24th

User Generated Videos

- Professional Produced Videos
 - Netflix
 - Hulu
- User Generated Videos
 - YouTube, Youku, Tudou
 - Hundreds of millions of short video clips
 - Wide ranges of topics
- Growing user generated videos
 - Readily available device
 - Production cycle is short







Motivation

- User experience in watching videos is not satisfactory
 - Slow startup time
 - Many pauses during playback



Measuring User Experiences Watching YouTube

Video download traces from various environments

Environment	Location	Network Technology
${ m E1}$	University 1	Campus WLAN
E2	Company 1	DSL
E3	Home 1	DSL
E4	Apartment 1	Cable Internet
E5	Dormitory 1	Campus LAN
${ m E6}$	Dormitory 2	Campus LAN
${ m E7}$	Apartment 2	Cable Internet
$\mathrm{E8}$	Town Library	Wireless Network
E9	Coffee shop	Wireless Network
E10	University 2	Campus WLAN
E11	Home 2	DSL
E12	Hotel	Wireless Network

Likelihood of Experiencing Pauses

- 10 out of 12 environments contain playbacks with pauses
- 41 out of 117 playbacks (35%) contain pauses



Number of Pauses

• 31 out of 117 playouts (22.6%) contain more than 10 pauses



How to improve user experiences?



Video Prefetching Scheme



- Prefetching Agent (PA)
 - Select videos to be prefetched and retrieve their prefixes
 - Store prefixes of prefetched videos
 - At clients (PF-Client) or proxy (PF-Proxy)
- Predict videos that are most likely to be watched
 - PA determines videos to prefetch from incoming requests

How to select videos to prefetch?

- PA predicts a set of videos to be requested
- Two main sources of video requests
 - Search Result lists
 - Related Video lists
- Use top N videos from these lists
- Advantages
 - Simple
 - Require no additional data
 - Effectiveness?



Datasets for Evaluation

Traces of data traffic between a campus network and YouTube servers

Trace File	T1	T2	T3
Duration	$1 \mathrm{day}$	3 days	$7 \mathrm{~days}$
Start Date	20-Oct-09	8-Jan-10	28-Jan-10
# Request	$71,\!282$	$7,\!562$	$257,\!098$
# Unique Clients	$7,\!914$	607	$10,\!511$
# Unique Videos	$48,\!978$	$5,\!887$	$154,\!363$

 Retrieve Search Result lists and Related video lists via YouTube data API

How Often Users Click on Related Videos and Search Results?

- Determine the referrers of each video request in the traces
 - From URL patterns, e.g., feature=related, feature=channel
 - From inference: look at a browse session to infer requests from Search Result list
- Related Video lists and Search Results lists are the most frequently used referrers



Evaluation Methodology

- Issue the requests based on real user request traces
- Keep track of the videos in PA's storage
- Evaluation metric
 - Hit ratio: How many requests we can serve from the PA's storage?

Hit ratio = <u>Hit requests</u> All requests

Effectiveness of various scheme combinations



- Videos from a Related Video list of a user are watched by other users
- Best combination is using RV-N algorithm with PF-Proxy setting

Combining Caching with Prefetching



- Cache-and-Prefetch can reach up to 81% of hit ratio
- Improvement is smaller as N increases due to larger overlapping between prefetched videos and cached videos

Analyzing Hit Ratios

- Only half of the hit requests come from RV lists
- Requests from SR lists is a large portion of the hit requests especially in PF-Proxy setting
- Recommendation system is a good indicator of topic interest







Analyzing the High Hit Ratios

 RV lists overlap with the video requests generated from other sources (esp. in PF-Proxy) up to 70%



Storage Requirement



- Measured in slots a slot holds one prefix of a video
- One slot = 2.5 MB (for prefix size of 30% and average video size of 8.4 MB)
- Require only 5 TB to reach 81% of hit ratio (at N=25)

Impact of Storage space



- Hit ratio decreases with the storage space size
- Still can achieve hit ratio of around 60% with 125 GB (50k slots)
- Compared to caching, cache-and-prefetch always performs better

Do we need to prefetch the whole video?



- Prefetching the whole videos is not necessary
- From analysis of video download traces, each location and each video requires different prefix size

Feasibility – Traffic Overhead

• Suppose prefix size = 15%, N = 11 and caching whole videos

Scheme	Hit Ratio	Normalized load
No scheme	0%	1.00
Cache-only	40%	0.60
Prefetch-only	66%	1.44
Cache-and-Prefetch	74%	1.02

- Caching helps reduce the traffic
- Pure prefetching yields higher hit ratio while increase traffic by 44%
- Combining the two results in highest hit ratio and only introduce 2% additional traffic

Conclusion

- Watching videos with prefix prefetching
 - Delay and Pauses are often
 - Prefix prefetching is feasible during browsing
 - Related videos are good interest predictors
 - Prefetching can reach hit ratio over 81% while caching can reach hit ratio of 40%

Q&A