Cloud Transcoder:
Bridging the Format and Resolution Gap between Internet Videos and Mobile Devices

Zhenhua Li, *Peking University*

Yan Huang, Gang Liu, Fuchen Wang, *Tencent Research*

Zhi-Li Zhang, *University of Minnesota*

Yafei Dai, *Peking University*

About Mobile Devices

- Mobile devices

  - more and more popular

  - more than PCs

iPhone
iPad
Android
Ultrabook
laptop
…

- Mobile traffic

  - only iPad accounts to 10% Internet traffic!

  - mostly headed for video streaming

YouTube NETFLIX hulu YOUKU 优酷 PPTV

“Gap” Between Mobile and Videos

- Today’s mobile video streaming is still challenging for a number of reasons
  - small and diverse screens

  - low battery power

  - embedded CPU

- Today’s Internet videos
  - mostly PC-oriented
  - single format (soft encode)
  - very limited resolutions

Format and resolution

Gap

Local transcoding

- Computation complexity of video transcoding

  - usually as 5 – 20 times as that of video decoding (viewing)

  - easily consume up the battery power of a mobile device

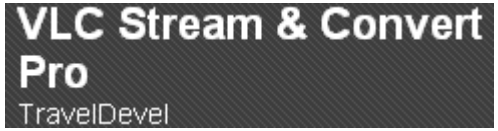

- So, today's mobile users often have to utilize their PCs with auxiliary software

  - iTunes, AirVideo, etc.

  - very inconvenient

Cloud-based transcoding

- Recent years, a worldwide upsurge of cloud service deployments

  - gradually move computation-intensive works from light-weight users onto heavy-weight clouds

- Traditional cloud transcoding solution

  - typically let users upload their original videos
  - work well for transcoding audios and short videos

  - unfit for long videos: 1. asymmetric Internet access (like ADSL), 2. Long videos consume very much computing resource, users need to wait a long time

youconvertit Beta

ONLINE-CONVERT.COM

movavi online convert

Multi-format support mobile player

- support full format video (373)
- Cannot not solve the resolution adaption problem
- http://player.qq.com

Cloud Transcoder [TENCNET]

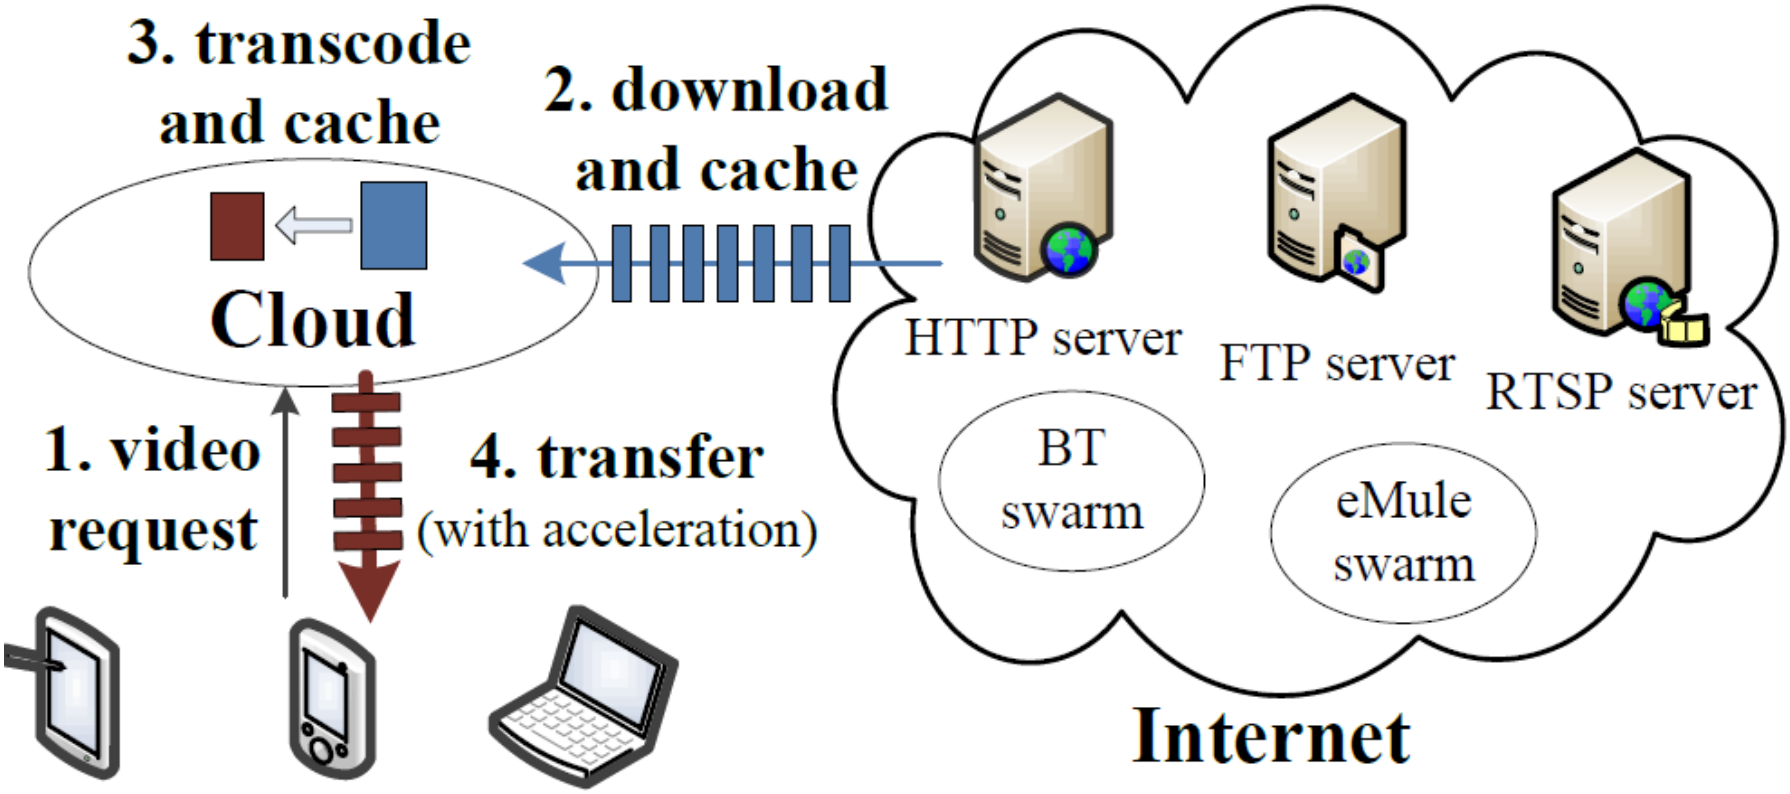

Looks simple and straightforward, while works effectively!

Work flow

- 1. The user only uploads a video request

  < video link; format, resolution, …>

HTTP/FTP/RTSP link
BT/eMule/Magnet link

User-specified transcoding parameters

- 2. The cloud caches both original videos and transcoded videos

- 3. The cloud transfers transcoded videos back to users with a high data rate

  - via the intra-cloud data transfer accelerations

  - **detailed described in Cloud Download Paper 2011 ACM MM**

Advantages

- Time Saver
  - Uploading time
  - Transcoding time

- Energy
  - Mobile user only consumes energy in the last step

    fast retrieving the transcoded video from the cloud

  - Cloud Transcoder provides **energy-efficient on-demand** video transcoding service to mobile users

Problem and solutions

- Cloud Transcoder moves all the video download and transcoding works from its users to the cloud

- So, a critical problem: how to handle the resulting heavy *download bandwidth pressure and transcoding computation pressure on the cloud*

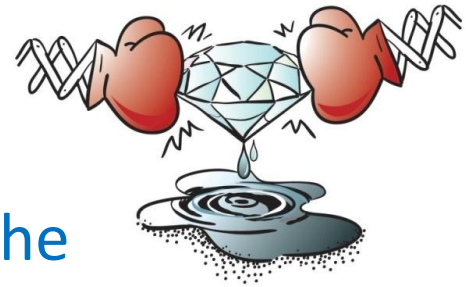

- Our solutions:

  - implicit data reuse among users via cloud cache

  - explicit transcoding recommendation and prediction

  - simple but effective: (1) download task cache hit ratio → 87%, (2) transcode task cache hit ratio → 66%

Real-world system

- Cloud Transcoder

  - deployed since May 2011
  - employs 244 commodity servers
  - across ten biggest ISP networks in China
  - serving ~8600 requests from ~4000 users per day
  - 96% original videos are long videos (> 100 MB)

  - system architecture is planned to serve 100,000 requests per day

Building Block	Number of servers	CPU (4 cores)	Memory	Storage	Bandwidth
ISP Proxy	6	Intel Xeon X3430 @2.4 GHz	8 GB	250 GB	1 Gbps (Intranet), 0.3 Gbps (Internet)
Task Manager	4	Intel Xeon X3210 @2.13 GHz	8 GB	250 GB	1 Gbps (Intranet)
Task Dispatcher	3	Intel Xeon X3210 @2.13 GHz	8 GB	460 GB	1 Gbps (Intranet)
Downloaders	20	Intel Xeon X3430 @2.4 GHz	8 GB	460 GB	1 Gbps (Intranet), ~0.325 Gbps (Internet)
Transcoders	15	Intel Xeon X3430 @2.4 GHz	8 GB	460 GB	1 Gbps (Intranet)
Cloud Cache	170 chunk servers, 23 upload servers, and 3 index servers	Intel Xeon 5130 @2.0 GHz	8 GB	4 TB (chunk server), 250 GB (upload server)	1 Gbps (Intranet), ~0.3 Gbps (Internet)

System Overview

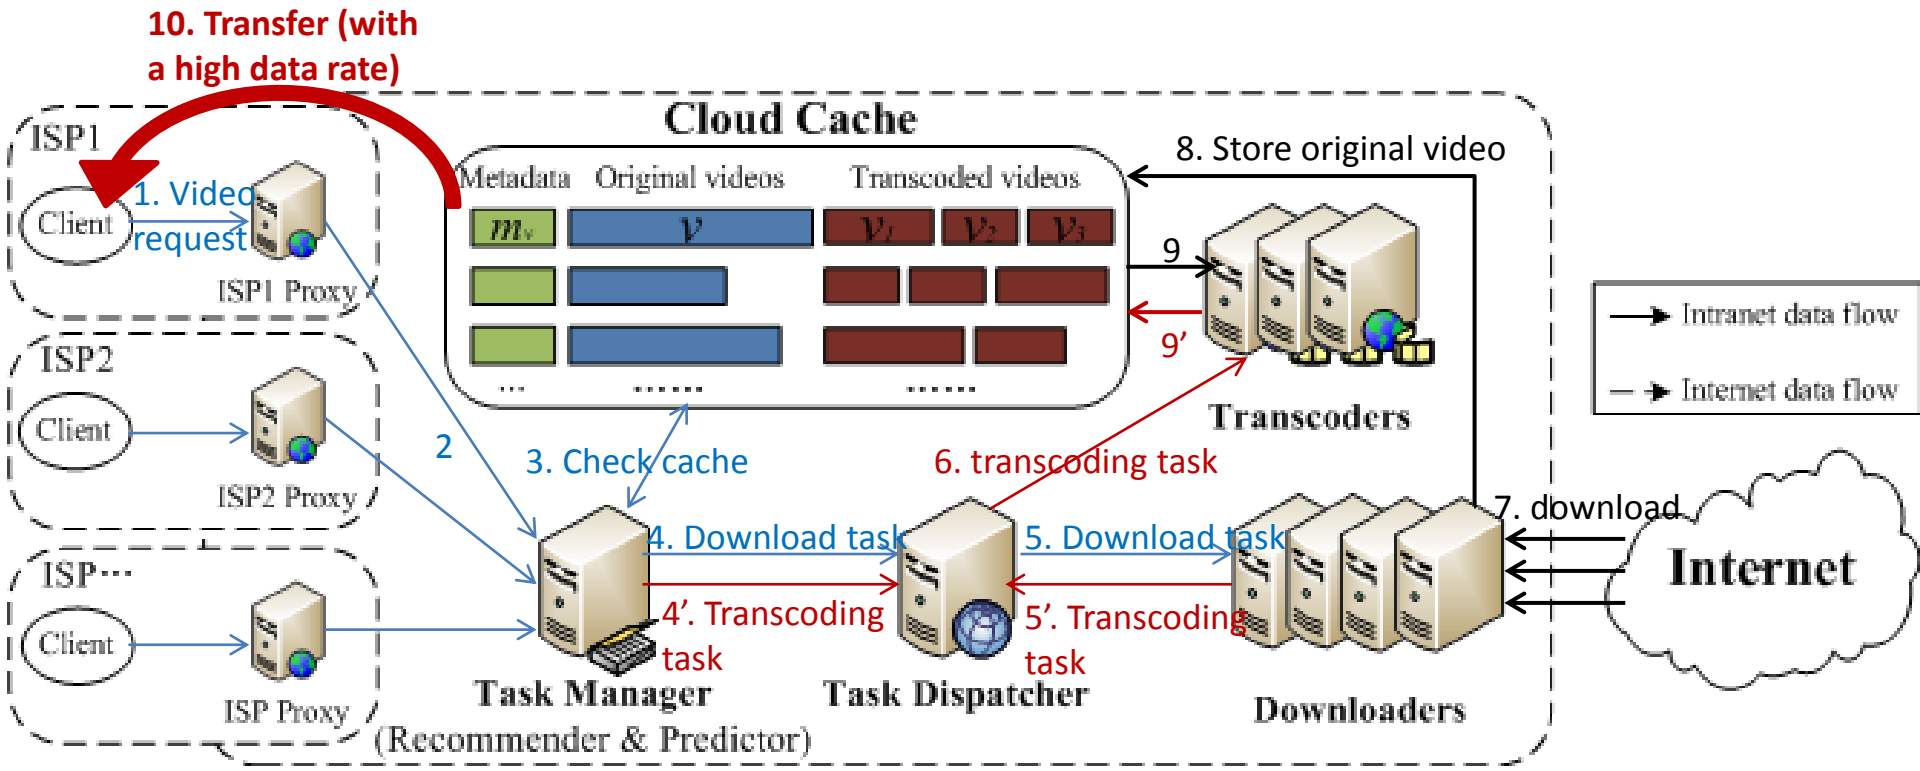

Transcoding Prediction

- When the average *computation pressure (CPU utilization) of the transcoders stays* below a certain threshold (50%) during a certain period (one hour)

  - Task Manager starts to predict which videos are likely to be requested for transcoding into which formats and resolutions

  - based on the video popularity information

  - Task Manager picks top-1000 popular videos and top-3 popular transcoding parameters to initiate transcoding tasks

  - part of the transcoding computation pressure in “hot” time has been moved to “cold” time for load balancing

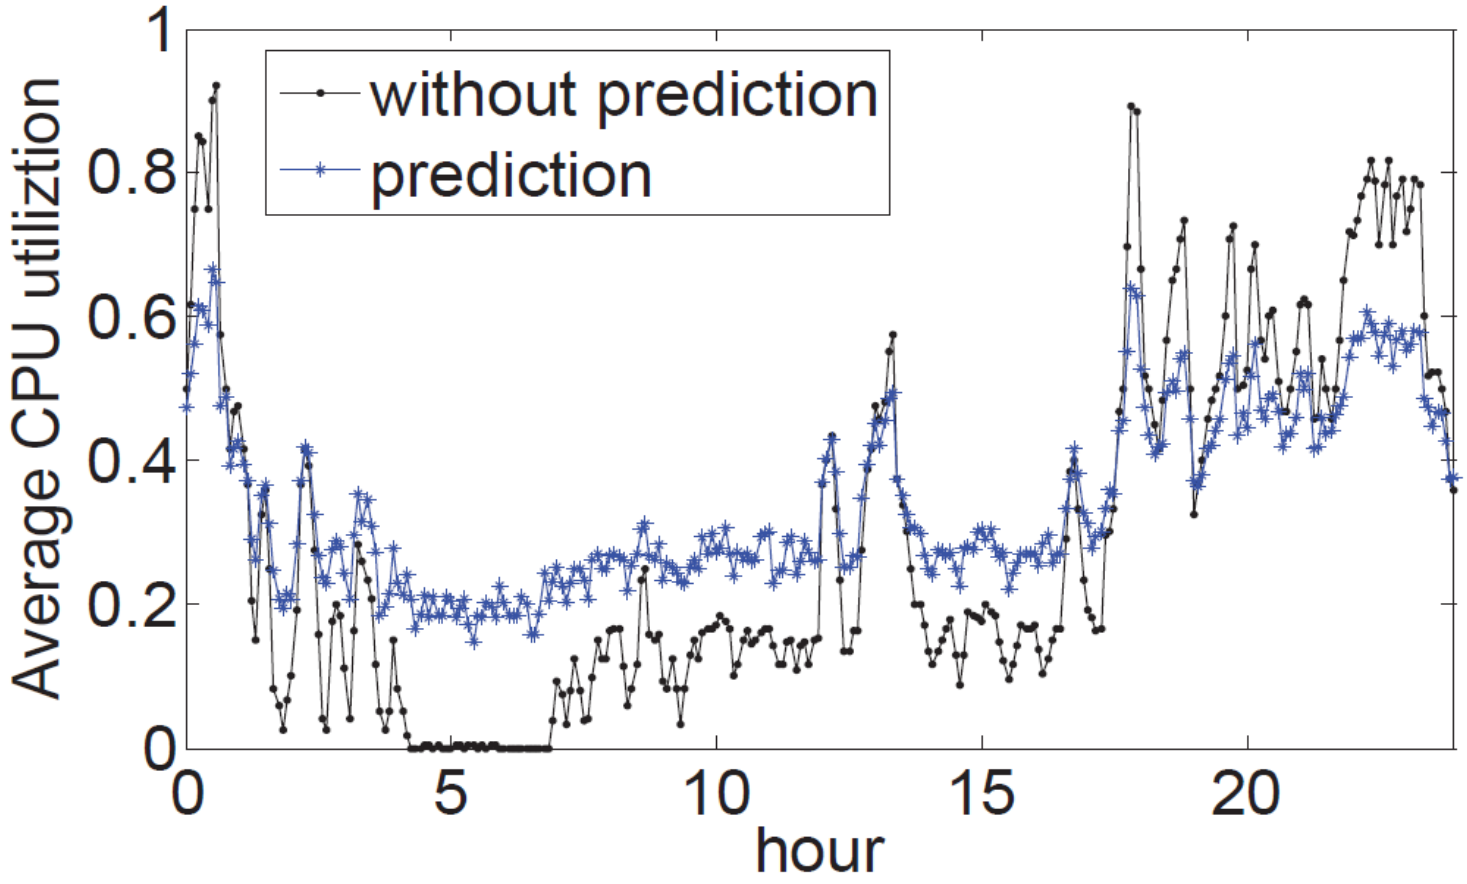

Figure 3: Average CPU utilization of the transcoders in one day (with prediction) and the other day (without prediction), respectively.

Cloud Cache

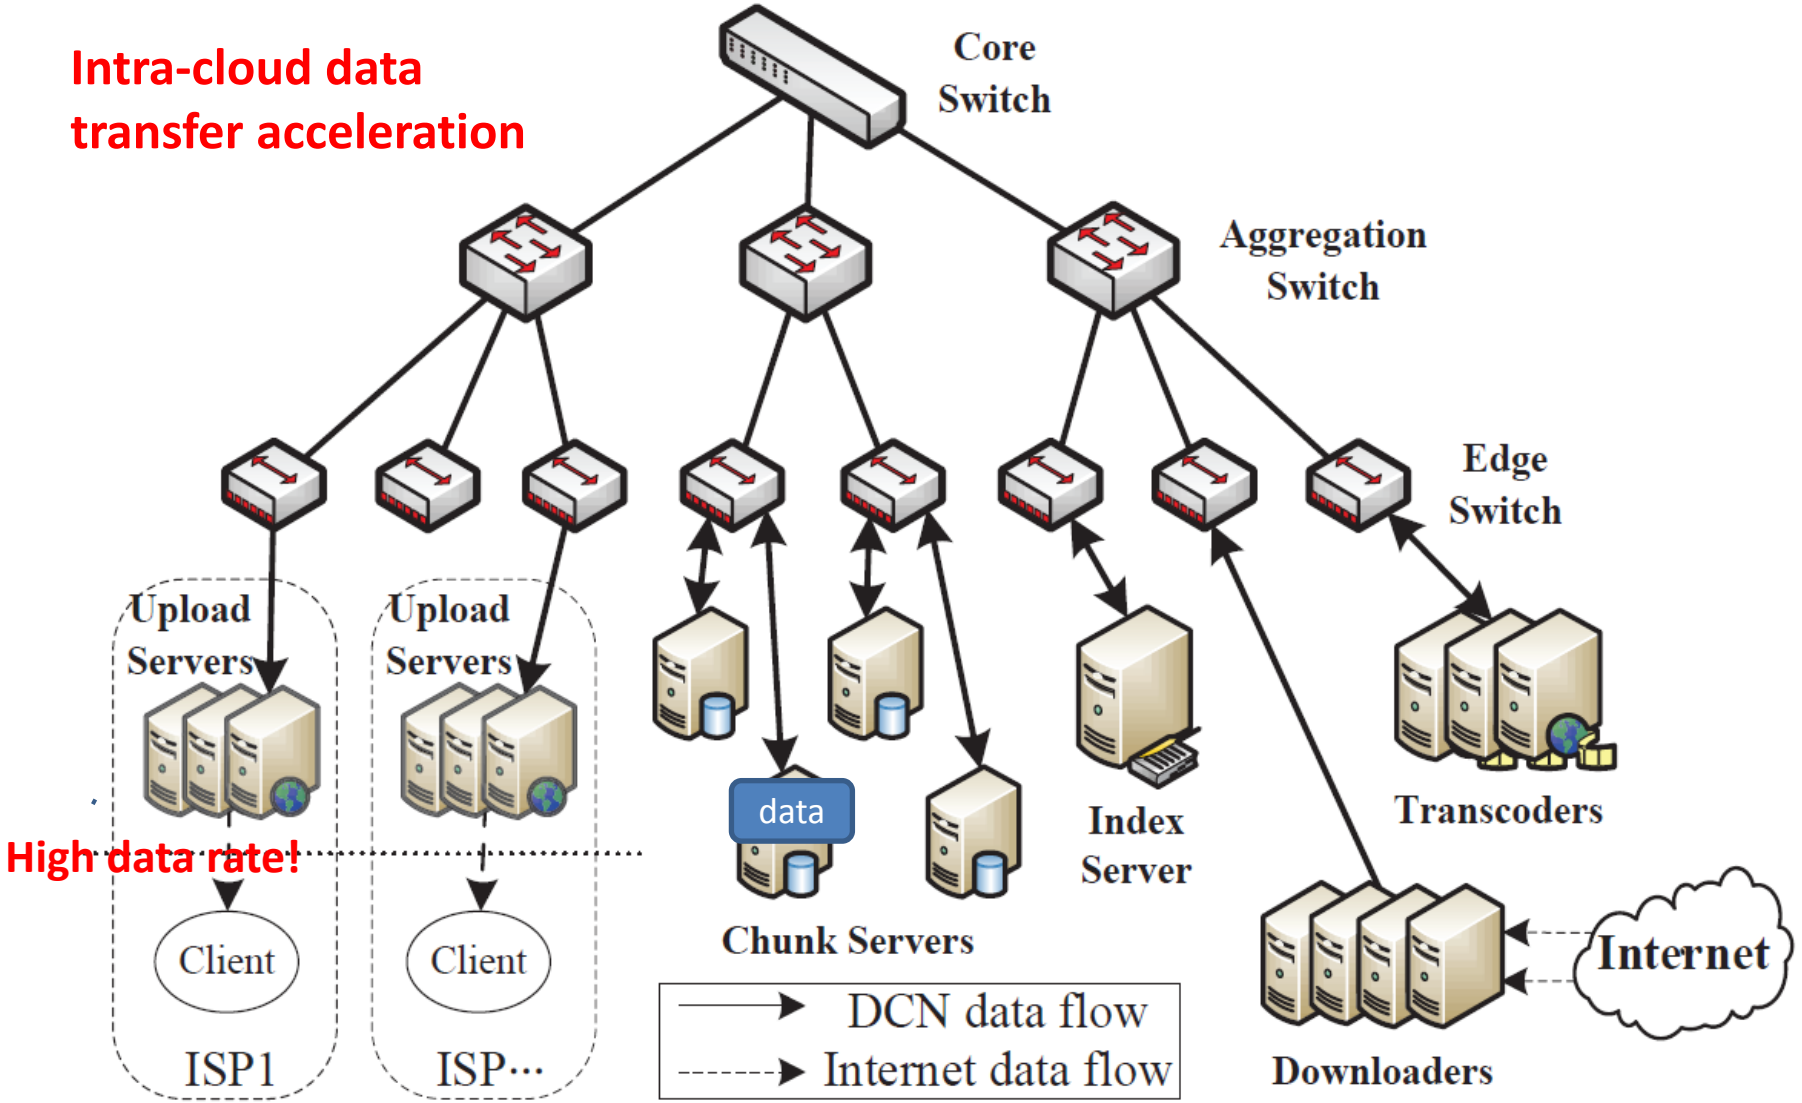
Intra-cloud data transfer acceleration
Core Switch
Aggregation Switch
Edge Switch
Upload Servers
Upload Servers
data
Index Server
Transcoders
High data rate!
Client
Client
ISP1
ISP···
Chunk Servers
DCN data flow
Internet data flow
Internet
Downloaders

Cloud Cache capacity planning

- Plan to handle 100K daily requests

  - avg size of original videos: 827 MB
  - a novel video is stored for 12 days
  - avg cache hit ratio of original videos: 87%
  - Original video cache capacity: C1 = 827 MB * 100K * 12 * (1-87%) = 126 TB

  - an original video has 3 transcoded videos in average
  - avg size of transcoded videos: 466 MB
  - Transcoded video cache capacity: C2 = 466 MB * 100K * 12 * (1-87%) = 213 TB

  - In total, C = C1 + C2 ≈ 340 TB

Cloud Cache replacement strategy

- Trace-driven simulations
- Compare FIFO, LRU and LFU
- LFU performs the best!

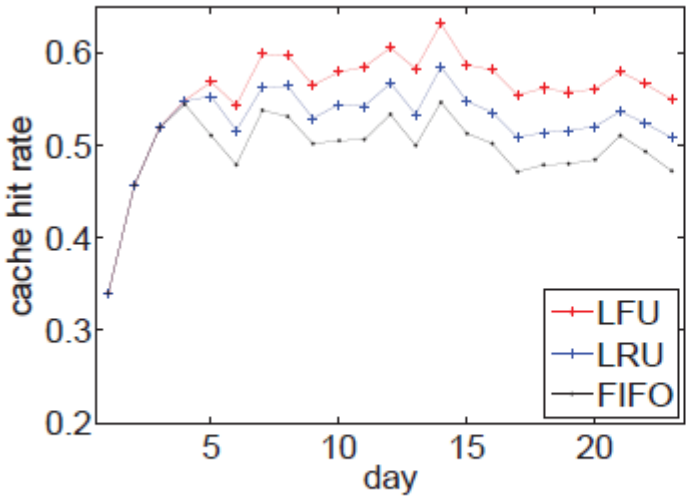

(a) For original videos.

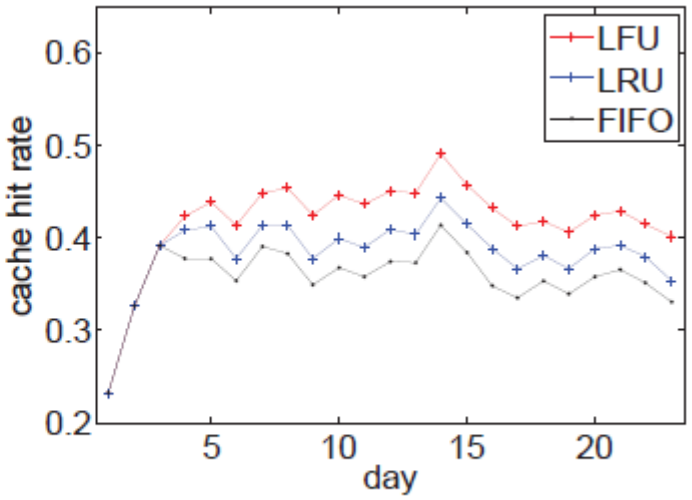

(b) For transcoded videos.

Performance Evaluation

- complete running log of Cloud Transcoder in 23 days (Oct. 1–23, 2011)
  - 197,400 video transcoding tasks involving 76,293 unique videos
  - 85% video links are P2P links
  - most popular transcoding parameters: (1) MP4-1024*768 (iPad), (2) MP4-640*480 (iPhone & Android), (3) 3GP-352*288 (Android)

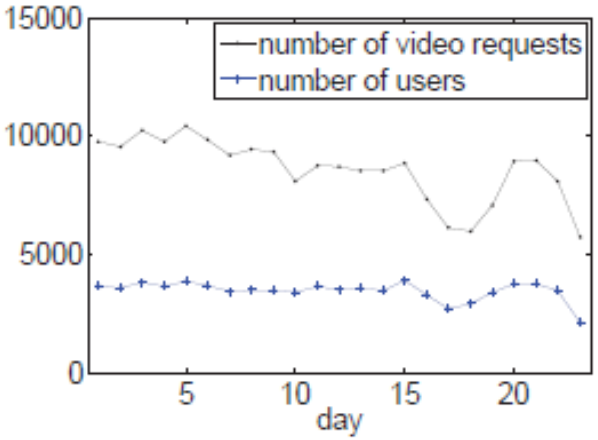

Figure 6: Daily statistics.

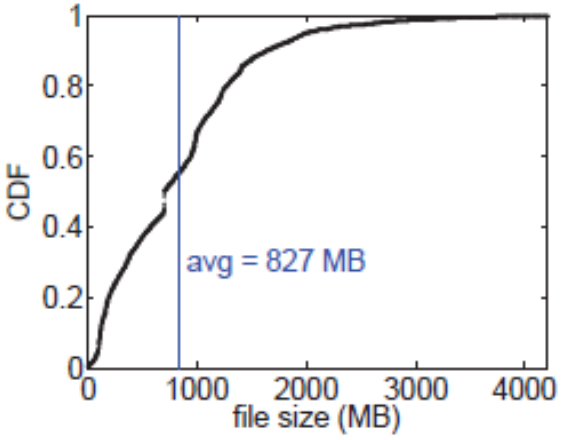

Figure 7: Original file size.

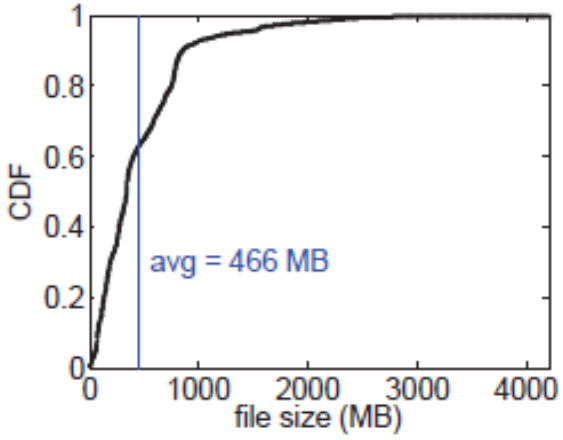

Figure 8: Transcoded file size.

Results

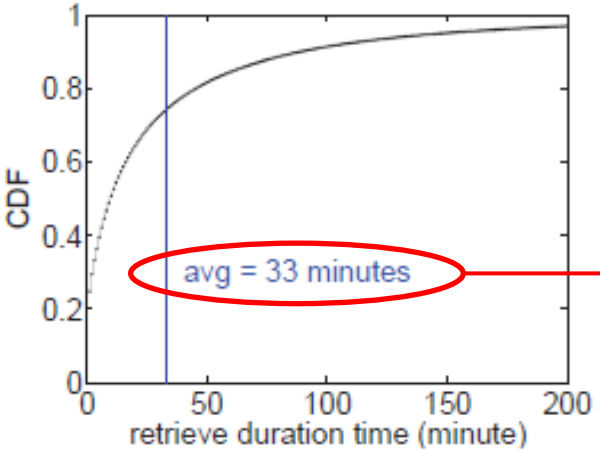

Figure 9: Retrieve duration.

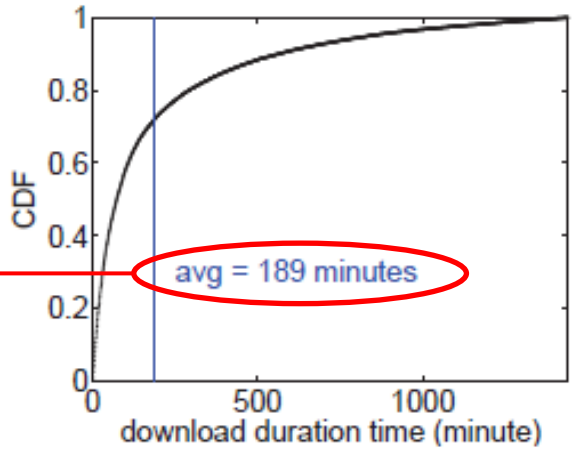

Figure 10: Download duration.

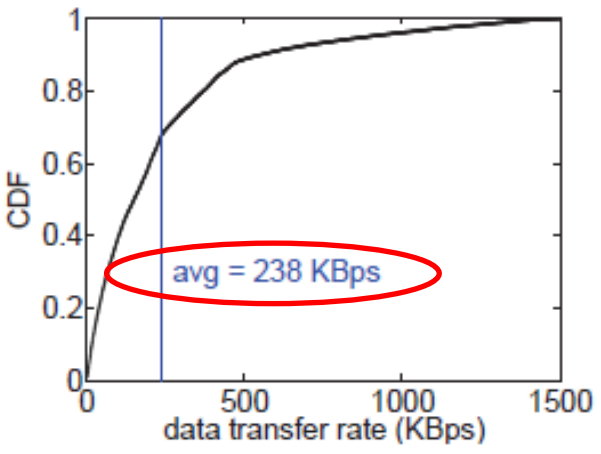

Figure 11: Data transfer rate.

Data transfer rate (≈KBps)	50	100	200	300
iPhone battery consumption (%)	8.7	8.9	9.0	9.2
iPad2 battery consumption (%)	4.5	4.8	5.0	5.1

Future work

- Cloud Transcoder: a novel prototype system
  - still at its startup stage
  - tend to adopt straightforward and solid designs
  - still considerable optimization space

- Other cloud transcoding services
  - mobile web browsers

Q & A