Measuring YouTube Quality of Experience for Users in Residential ISPs*

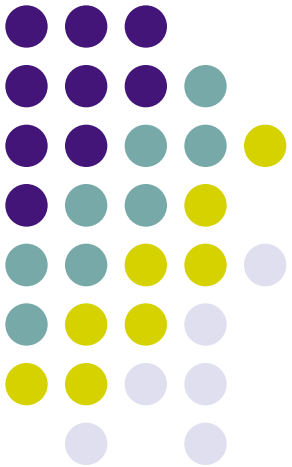

Deep Medhi

dmedhi@umkc.edu

University of Missouri–Kansas City

February 2013

**Joint work with Parikshit Juluri (UMKC), Louis Plissonneau (Orange Labs, France), and Yong Zeng (UMKC)*

QoE with YouTube

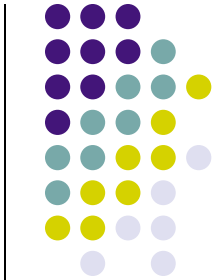

- Need a measurement-based approach
- Passive vs Active approach
- Passive: collect trace data at some ‘collection point’
  - But cannot quite get end users’ QoE
- Active:
  - Subjective studies with users accessing and collecting data
    - Users need to access video and data needs to be collected – collection is not precise, besides difficulty
  - Bot-based approach: possible from Campus Network or PlanetLab, but most such networks have high-speed access, so not really QoE of actual end users from ISPs

Broad Goals

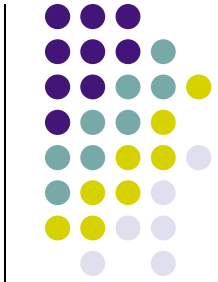

- Collect QoE metrics for YouTube video streaming sessions from end users
  - Important to collect the metrics precisely
- Compare the QoE of end-users across multiple ISPs
- Additionally address, how does YouTube's CDN impact users in different ISPs?

- *With the above goals: Pytomo was developed!*

Pytomo

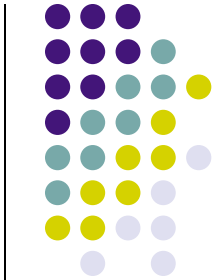

- Active measurement tool developed in Python: “emulate a video player in a browser”
- Emulates users’ YouTube viewing experience in a browser
  - collects data precisely (such as duration of interruptions etc)
- It actively crawls the YouTube website to progressively download a number of videos (i.e., the user does not have to launch the tool for each video)
- Can collect CDN information, Latency Metrics and also QoE metrics.
- Open source: runs on multiple platforms (Linux, Mac, and Windows) with minimal initial setup.
  - Available for download at : http://code.google.com/p/pytomo/

Running the Tool

- For Windows (most residential users use Windows):
  - Download the exe file and run it by entering a few basic pieces of information such as how many videos to crawl. It's that simple!!!

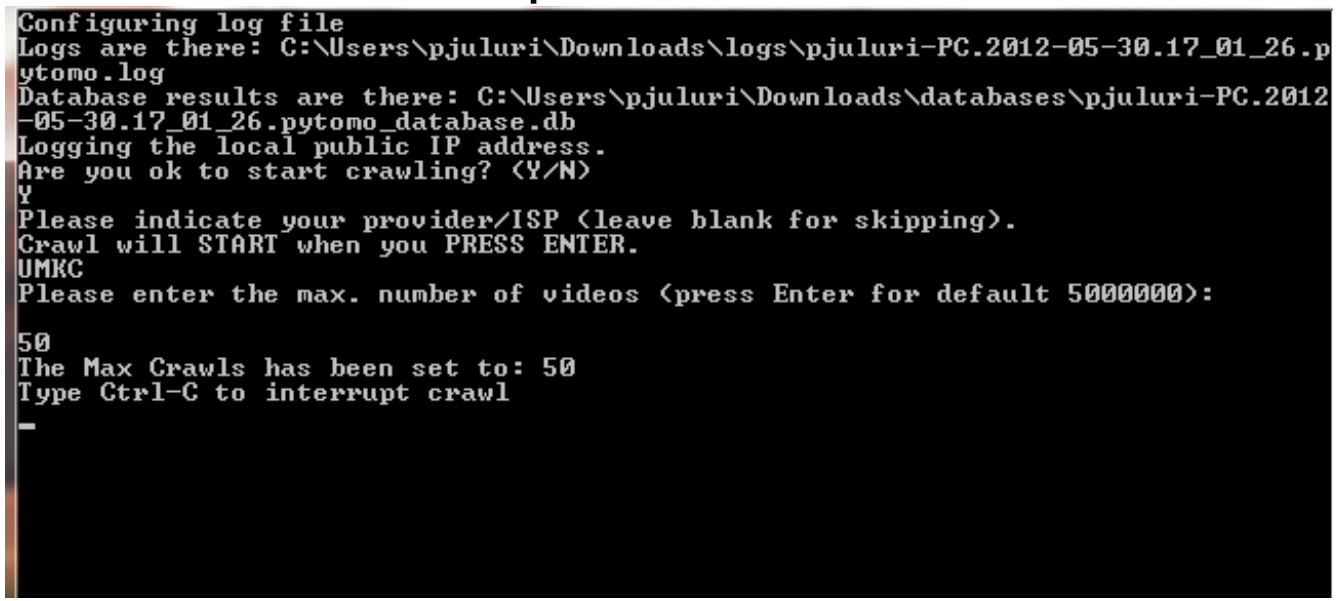

Pytomo: Steps To Collect Data

1. **Select initial set of videos:** Select an initial set of videos (VideoList) to start the crawl. It selects the most popular videos world-wide within a configurable time-frame (day, week, month, or all-time) at the time of accessing
2. **Obtain URL of the server hosting the video:** From the HTML page of the video the URL of the video-server hosting the video file(FLV, mp4 etc.,) are obtained and the IP address of the URL is resolved.
3. **Obtain latency metrics:** The latency towards the video server is measured by collecting ping metrics

Pytomo: Steps To Collect Data

- **Obtain meta-data of video:** While downloading the video, collect meta-data such as the VideoDuration, VideoLength, and EncodingRate.

- **Obtain playback statistics of video:**
  - Emulation mechanism: emulate a video viewing session as perceived by a user. The emulation is achieved by using two different time trackers.
  - **DownloadTimeTracker:** Keeps track of the download of the video and follows the absolute download time.
  - **PlaybackTimeTracker:** Keeps track of the playback time of the video, thus emulating the progressive download. PlaybackTimeTracker is incremented according to the video tags found in the video.

Parsing With FLV

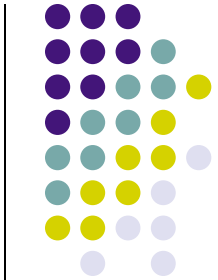

- Flash video (FLV) is an encoding format
- An FLV consists of a short header, and then an interleaved stream of audio, video and metadata packets (tags).
- Each tag has the following fields:
  - Type : Audio, Video or Meta
  - Body Length : Size of the body
  - Timestamp : Timestamp of tag in milliseconds
  - Media Payload
  - Previous Tag Size : Total Size of previous tag, or 0 for first tag
- While parsing the FLV file our tool reads the timestamp on each tag to determine the precise time when it will be played, thus ensuring that being able to precisely emulate the FLV video playback.

Pytomo Steps

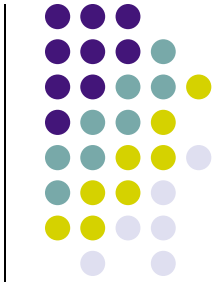

The playback of the video goes through the following stages:

1. **Initial buffering stage:**
   - Time for which the video is buffered before playback begins.
   - The Initial Buffer Size is a measure of amount of data required to play 2 sec of a video (was determined through controlled study)
   - This buffer size value is stored as InitialData and the data rate at which this buffer is filled is saved as Initial-Rate.
   - From InitialData and InitialRate, we can calculate StartTime.
2. **Playback stage:**
   - The state where the initial buffer is filled and the video playback begins.
   - In case the number of bytes of a video in the buffer falls below the Buffering Threshold, we interpret this as an interruption in the playback and the PlaybackTimeTracker is paused. The number of times a video is interrupted during playback is tracked as DownloadInterruptions

View of The Entire Process

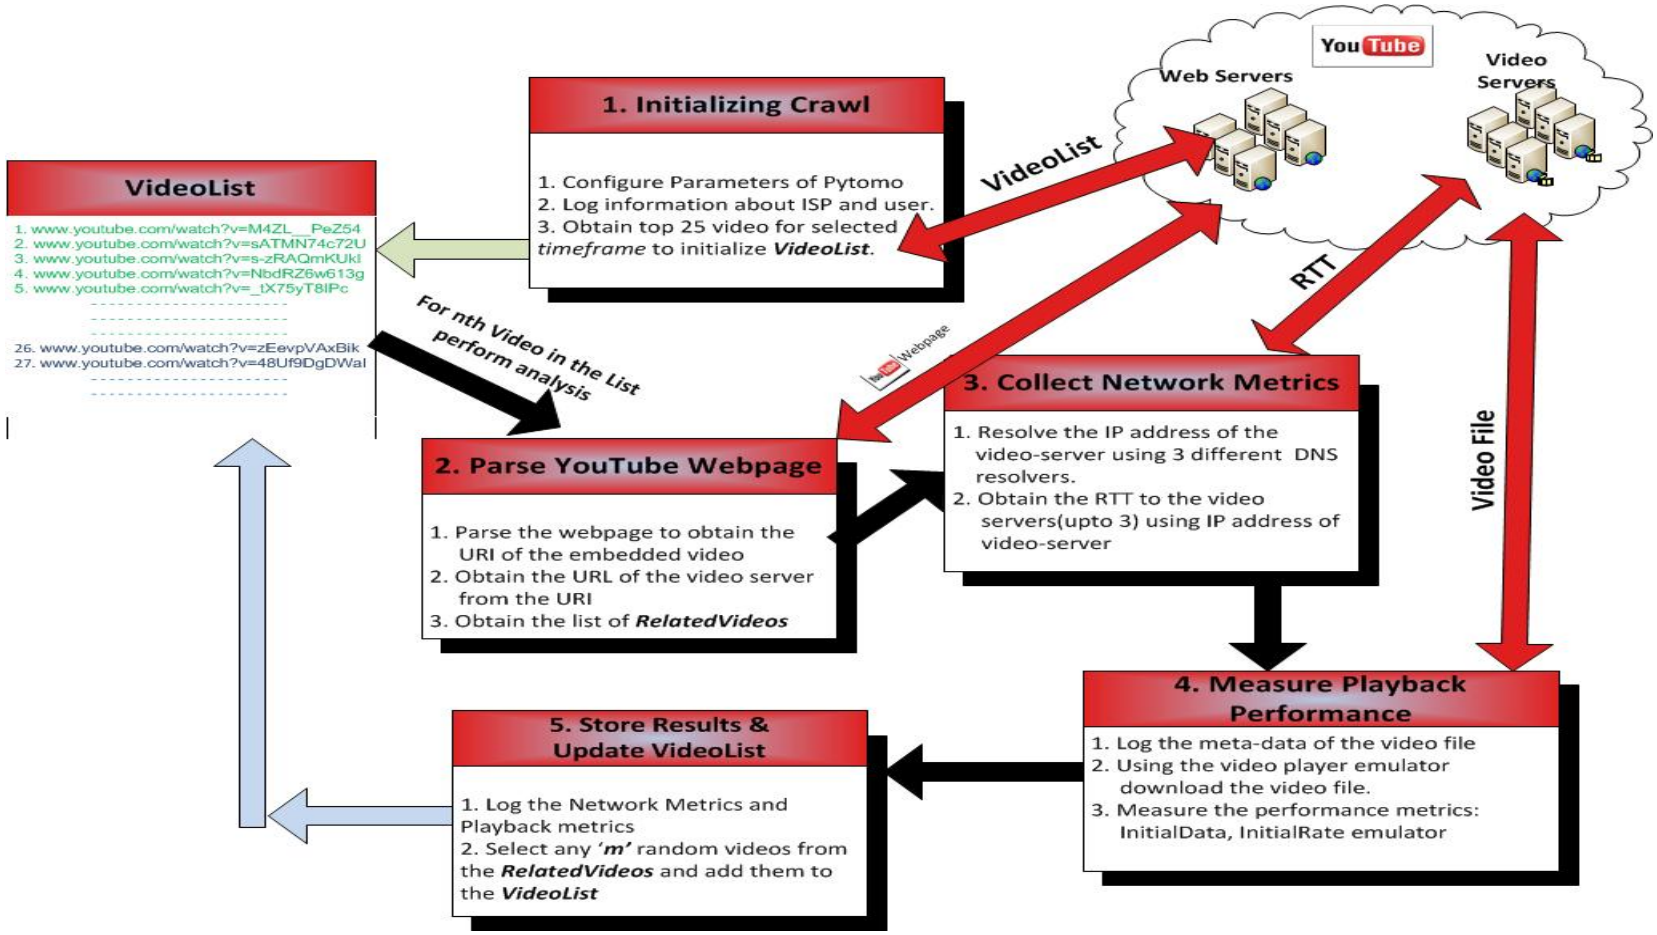

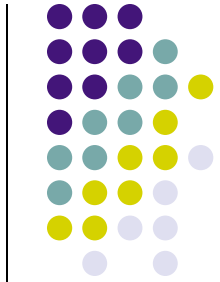

Next:

Understanding QoE through a Measurement Study

Data Collection in Kansas City

- Why? Two songs☺
  - "*Everything's up to date in Kansas City*": from musical "Oklahoma" Rodgers & Hammerstein

    http://www.youtube.com/watch?v=B_x_67fXtqM
  - "*I'm goin' to Kansas City, Kansas City here I come*" - Jerry Leiber and Mike Stoller's song (1952)
  - http://www.youtube.com/watch?v=mypHZmXdU3o

In other words,

- Convenient☺
- But where is Kansas City?

Kansas City

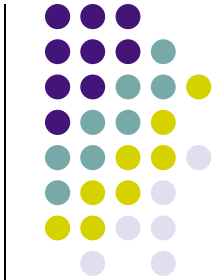

- It’s located centrally in the US
- Note that most large carriers’ fiber optic networks go through Kansas City

Data Collection in Kansas City

- This has been still a major challenge!!
- Wanted to do a controlled study
- Difficulty Identifying volunteers located throughout the Kansas City metropolitan area to run the Pytomo tool on their local machines through their residential ISPs.
- Volunteers were students at UMKC who were taking a course on Statistics
- Crawls were performed between 8:00 p.m. and 10:00 p.m. in the evening when residential use of the Internet is usually high.
- Two sets of crawls were performed:
  - December 2011 Crawls:
    - The 1st set of crawls were performed by 19 users on December 8, 2011.
    - A total of 1,289 videos were downloaded from 442 distinct video servers.
  - March 2012 Crawls:
    - The 2nd set of crawls were performed on March 14 and 15, 2012 by 32 different users.
    - A total of 2,390 videos from 988 distinct video servers were downloaded
  - Another recent collection (October 2012): to discuss later
- Users were located in three residential ISPs. (ISP1, ISP2 & ISP3)

Results

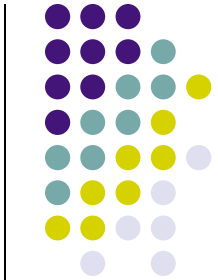

Considered data from two different aspects:

1. Video Server Selection Dynamics
   - How does the video server selection vary with ISP?
   - For the same ISP how does the video server selection vary between the two sets.
   - How does popularity of the video effect the server selection?
2. Analysis of Users' QoE
   - Compare the perceived QoE of users between different ISP's using QoE metrics

Video Server Selection Dynamics

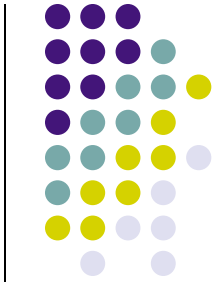

- Every YouTube video HTML page consists of a Java based YouTube player.
- The player obtains the video file to be played from a video server by using a URI that directs to the location of the video on a video server.
- For December 2011 and March 2012 measurements, we found that the video server URL is typically of the form:

  http://o-o.preferred.SERVER_CODE.v[1-24].lscache[1-8].c.youtube.com

  SERVER CODE identifies a server ID, behind which is likely a cluster of video servers in a data center environment.
- Typically, SERVER CODE indicates the geographic location of video server clusters.
- Commonly named by including IATA airport codes: For example:
  - dfw : Dallas-Forth Worth,
  - iad : Washington, DC
  - ord : Chicago

A pictorial view (Top 10 videos, December 2011 crawl)

- Left indicates users in a particular ISP and right indicates the server locations from which they are served

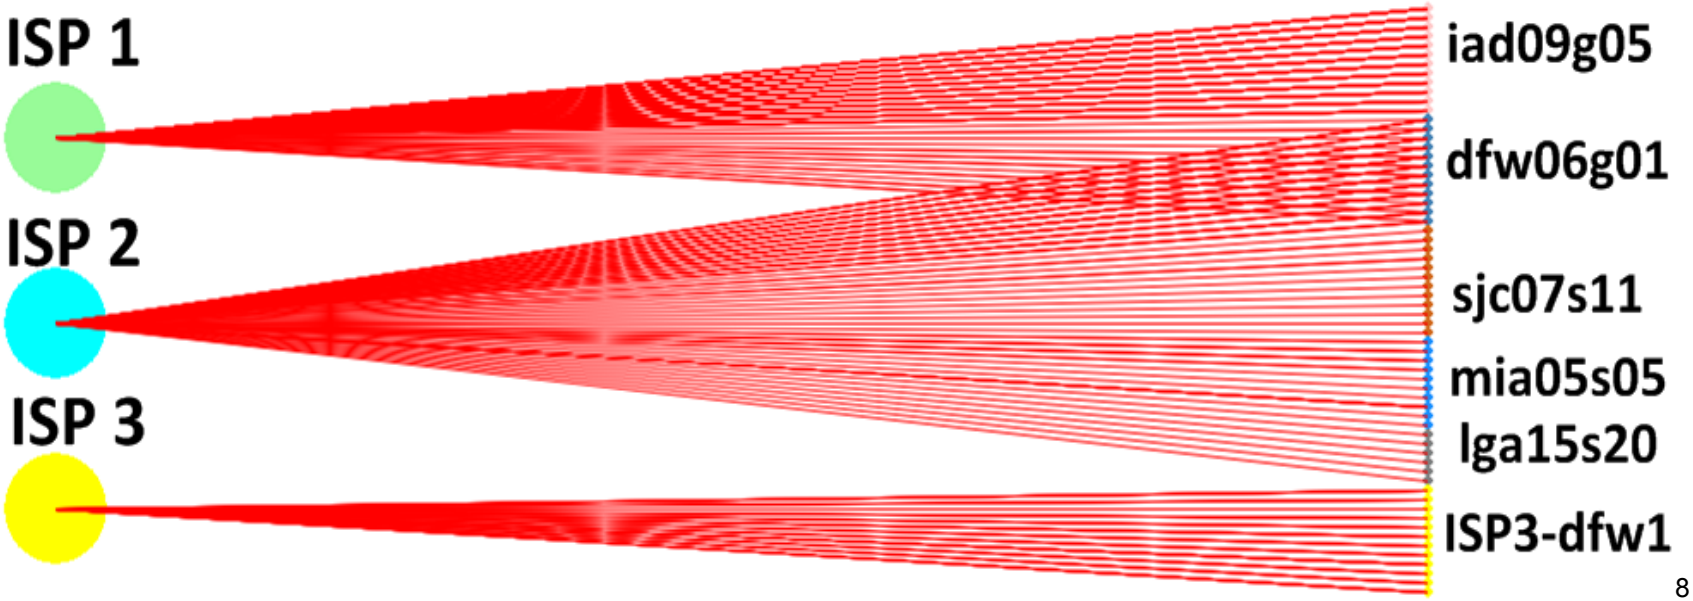

Server Selection for Top 10 Videos: December Crawl

- For most videos users of respective ISP's accessed exclusive videos
- Only Server ID *dfw06g01* was accessed by users in ISP1 and ISP2
- No other video servers were accessed by users of more than one ISP.
- ISP3 has access to exclusive servers for the top 10 videos.

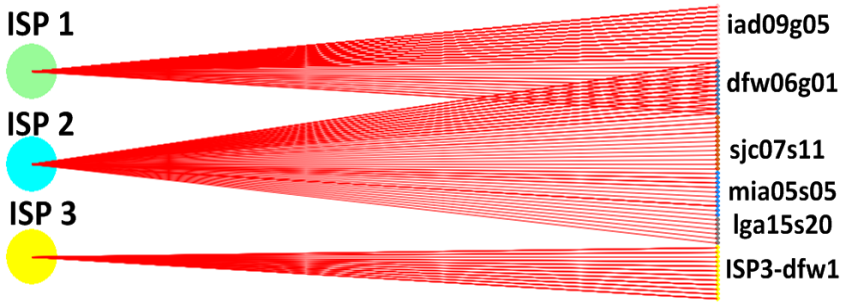

Server Selection for Top 10 Videos: March 2012 Crawl

- Similar Access pattern for the top 10 videos as in December.
- Server ID's *xo-ord1 & ord12s05* were accessed by users in ISP1 and ISP2
- Even in this crawl ISP3 has access to exclusive servers for the top 10 videos.

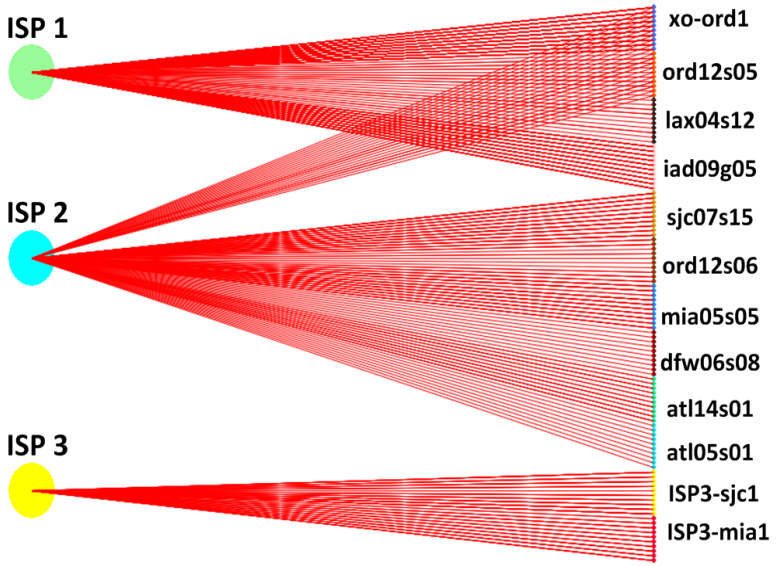

Server Selection for All Videos: March 2012 Crawl

- Access pattern follow the same trend even for the not so popular videos
- As seen for Top-10 videos in March, only *xo-ord1 & ord12s05* are common between ISP1 and ISP2
- ISP3 has access to exclusive servers even for non-popular videos
- ISP3 accesses servers in the geographic location as ISP2 (sjc & mia) but the clusters are different
- The presence of ISP-3's name in the SERVER CODE of the URL Supports the above statement

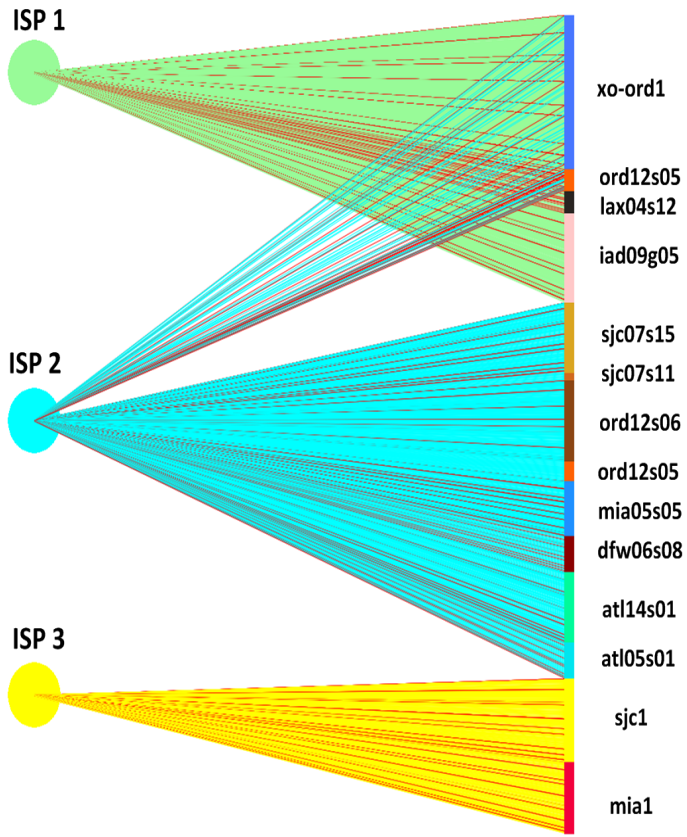

Change in Video Server Dynamics: from December 2011 to March 2012

- For ISP1, of two video server clusters that were observed in December 2011, only one was observed in March 2012 (PLUS three additional).
- For ISP-2, three out of the six observed in December 2011 were observed in March 2012 (PLUS six additional).
- Thus, it seems video servers allocated to the same ISP also change over a period of time, rather frequently.
- Users of ISP-3 were always served by dedicated clusters at different locations.

Analysis of Users' QoE

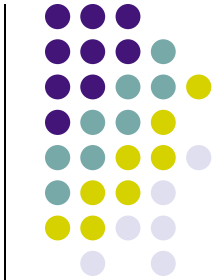

- Factors to compare the QoE of the users are:
  - *EncodingRate*
  - *StartTime*
  - *InitialRate*
  - *Average Download Rate*
  - *DownloadInterruptions*

Access Patterns with QoE

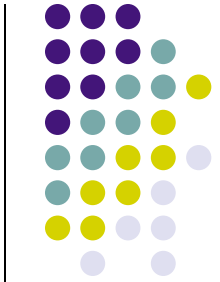

Table 1: Access Patterns for The Video-Server from 3 ISP's with QoE Metrics

Server Code	December 2011				March 2012			
	Samples	Avg. Start Time	% of Videos Interrupted	Avg. D/w (Mbps)	Samples	Avg. Start Time	% of Videos Interrupted	Avg. D/w (Mbps)
ISP-1								
all (ISP-1)	654	1.18	10.24	1.35	1468	0.49	1.90	2.60
dfw06g01	41	1.45	9.75	1.30	-	-		
iad09g05	613	1.15	10.29	1.35	213	0.65	2.81	1.82
xo-ord1	-	-	-	-	1213	0.44	0.99	2.82
lax04s12	-	-	-	-	20	0.97	0.00	1.69
ord12so5	-	-	-	-	22	1.44	45.45	0.55
ISP-2								
all (ISP-2)	347	0.77	1.80	1.50	602	1.23	16.00	1.78
dfw06g01	175	0.89	2.85	1.35	-	-	-	-
iad09s12	6	1.50	0.00	1.14	-	-	-	-
lga15s20	36	0.62	5.56	1.76	-	-	-	-
dfw06s10	3	0.85	0.00	1.58	-	-	-	-
dfw06s08	-	-	-	-	46	1.42	4.34	0.96
mia05s05	65	0.42	0.00	1.78	93	0.67	1.07	2.52
sjc07s11	62	0.83	0.00	1.50	6	2.71	33.33	1.06
sjc07s15	-	-	-	-	147	2.91	43.54	0.67
atl14s01	-	-	-	-	88	0.34	0.00	2.77
atl05s01	-	-	-	-	28	0.34	0.00	2.42
ord12s05	-	-	-	-	44	1.30	54.54	0.50
ord12s06	-	-	-	-	109	0.52	4.35	1.82
xo-ord1	-	-	-	-	41	0.34	0.00	3.86
ISP-3								
all (ISP-3)	259	0.78	0.59	1.68	320	0.63	0.63	1.84
ISP3-dfw1	259	0.78	0.59	1.68	-	-	-	-
ISP3-mia1	-		-	-	132	0.54	1.53	2.02
ISP3-sjc1	-		-	-	188	0.68	0.00	1.73

Additional Observations

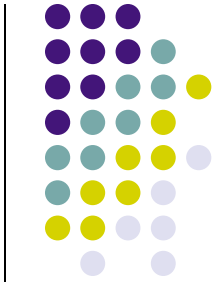

- Can we identify whether it's the network problem or server problem?
  - Two server IDs located in Chicago: ord12s05 and ord12s06; the former one is used by both ISP-1 and ISP-2, while the latter one is used only by ISP-2.
  - Since accessing ord12s05 shows high interruptions (and low download rate) for both providers, but low interruptions in case of ord12s06,
  - Very likely that ord12s05 is experiencing congestion

Encoding Rate

- The Encoding Rate determines the quality of the videos.
- The CDF indicates that the videos downloaded were of similar quality when accessed from different ISP's.
- Any variation in QoE is because of Network Performance and not the video format.

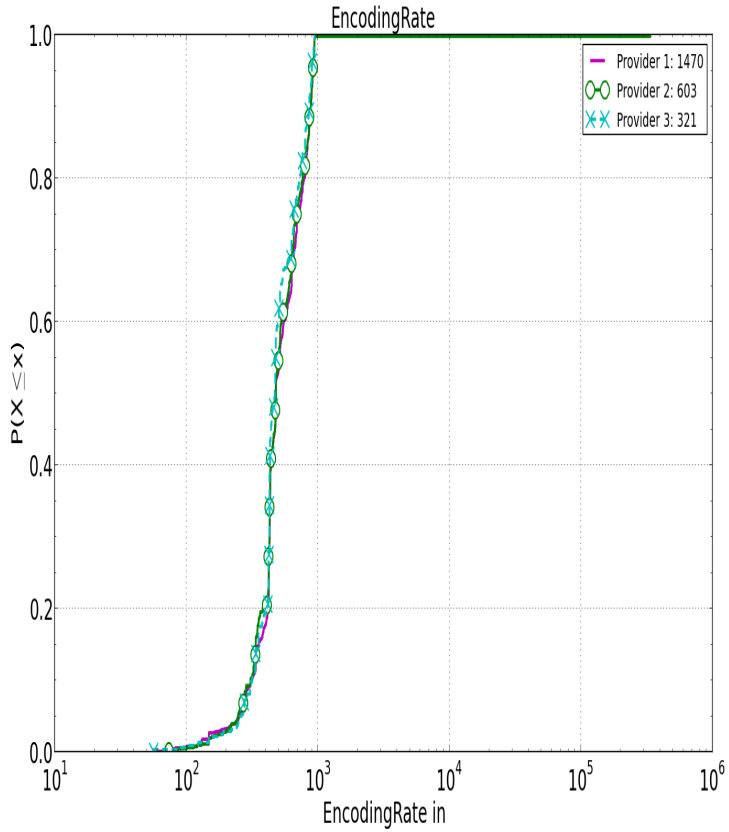

StartTime

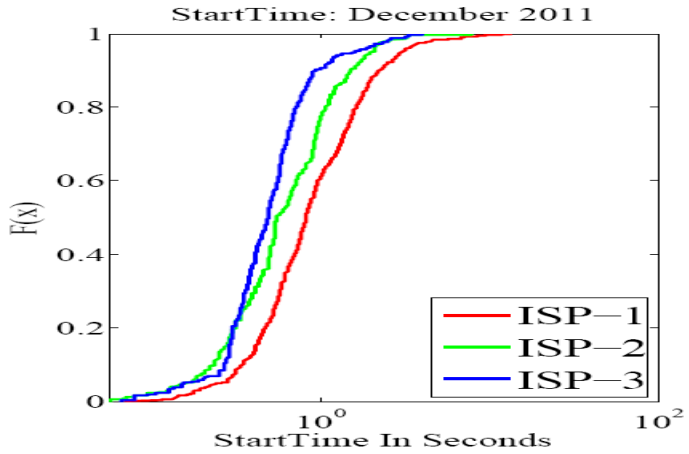

December 2011

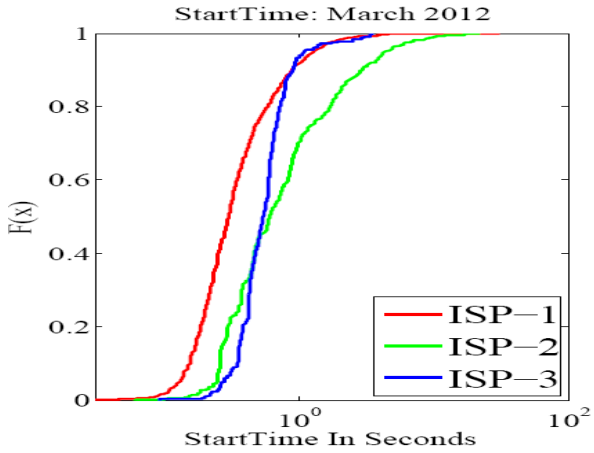

Match 2012

- **StartTime:** The amount of time the user has to wait before the playback begins.
- Average start time = less than 1 sec
- for ISP-2, at least two server locations experienced average start time of 2 sec or more in March 2012.

InitialRate

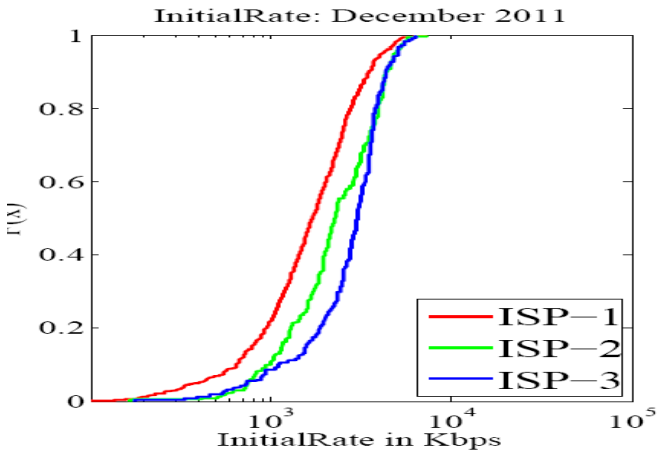

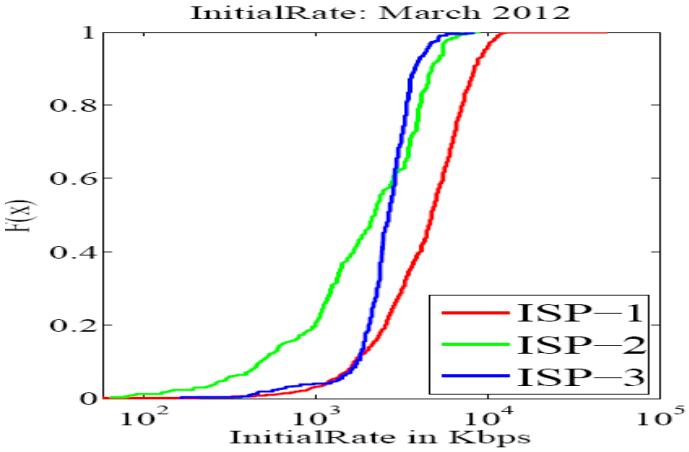

- Download rate observed during the initial buffering phase
- Significant increase in performance for ISP1
  - Change in the servers that were accessed.

Two Sample t-test (Dec vs Mar)

Table 2: p-value for two-sample t-test

p-value	ISP-1	ISP-2	ISP-3
DownloadRate	2.2×10^{-16}	0.52	0.22
StartTime	2.2×10^{-16}	1.18×10^{-8}	0.44

H_0: $\mu_{Dec} = \mu_{Mar}$ vs. H_{alt}: $\mu_{Dec} \neq \mu_{Mar}$

- For ISP-1 and ISP-2, StartTime in March 2012 is highly statistically significantly different than December 2011
- For ISP- 1, StartTime was less in March 2012, it was higher for ISP-2. There is no statistical difference for ISP-3.
- DownloadRate, for ISP-1, the March 2012 crawl is highly statistically significantly different (actually higher) than in the December 2011 crawl.
- No statistical difference for Download-Rate for the other two providers.

Download Interruptions

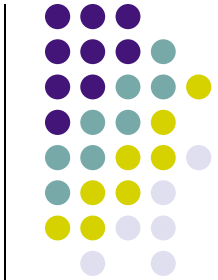

Table 3: Download Interruption Statistics

ISP	Number of videos		Videos w/ Interruptions(%)	
	December	March	December	March
ISP 1	654	1469	4.5	1.9
ISP 2	377	602	1.8	16
ISP 3	258	320	0.78	0.625

- For ISP-3: least number of interruptions

  For ISP-2, interruptions were rather high during in March 2012
- (Next table on servers that contribute to the maximum interruptions)
- New server ID sjc07s15 served the most number of videos in March for ISP-2 had an unusually high interruptions contributed to increase in the the rate to 16% (from 1.8%)
- Strong corelation between download rate and percentage of interruptions (not surprising)

Download Interruptions

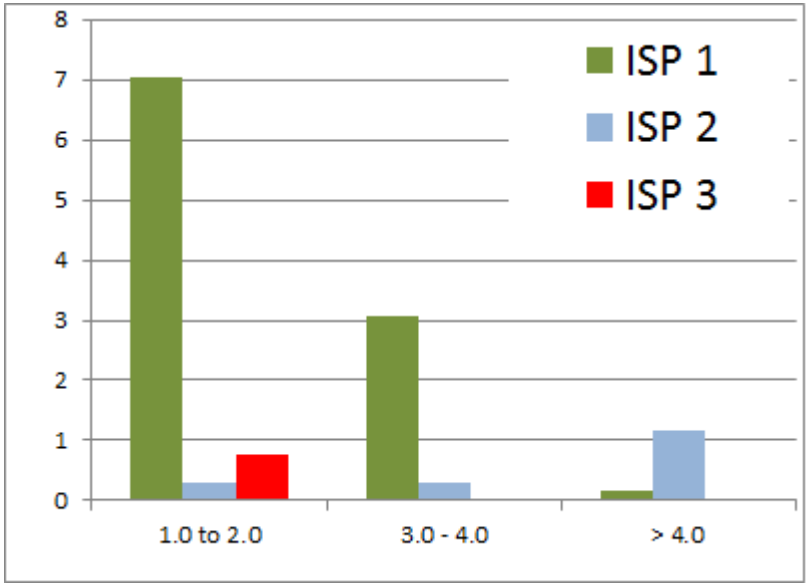

December 2011

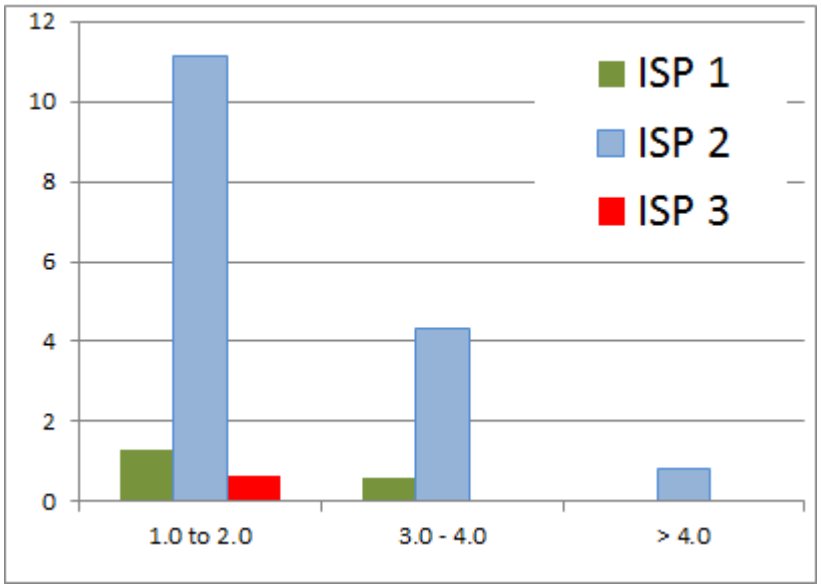

March 2012

Size of circles=number of videos served Color (avg. ping time), green, blue, red

- Green (< 50 ms)
- Blue (50 – 200 ms)
- Red (> 200 ms)

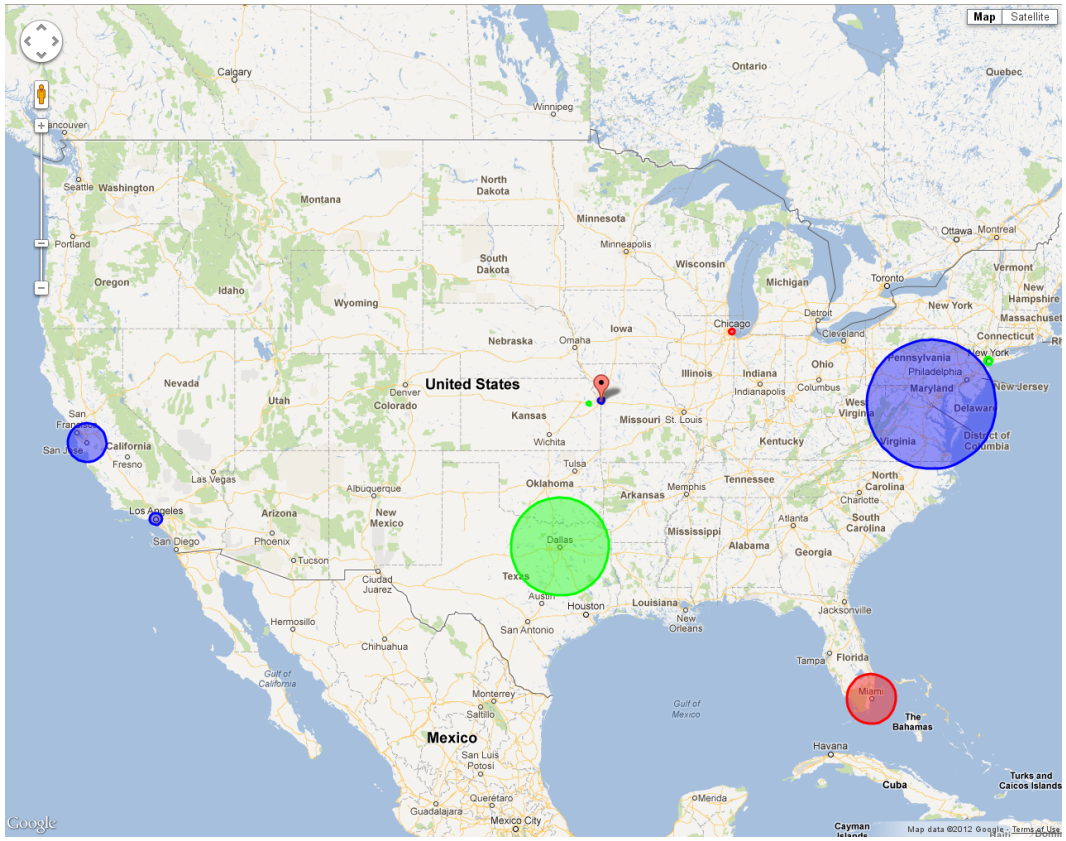

Recent Data Set (October 2012)

- Serving naming convention has changed!
  - No information about city code in the name☹
- We did some checking with previous data set and make inference about locations for some of them.

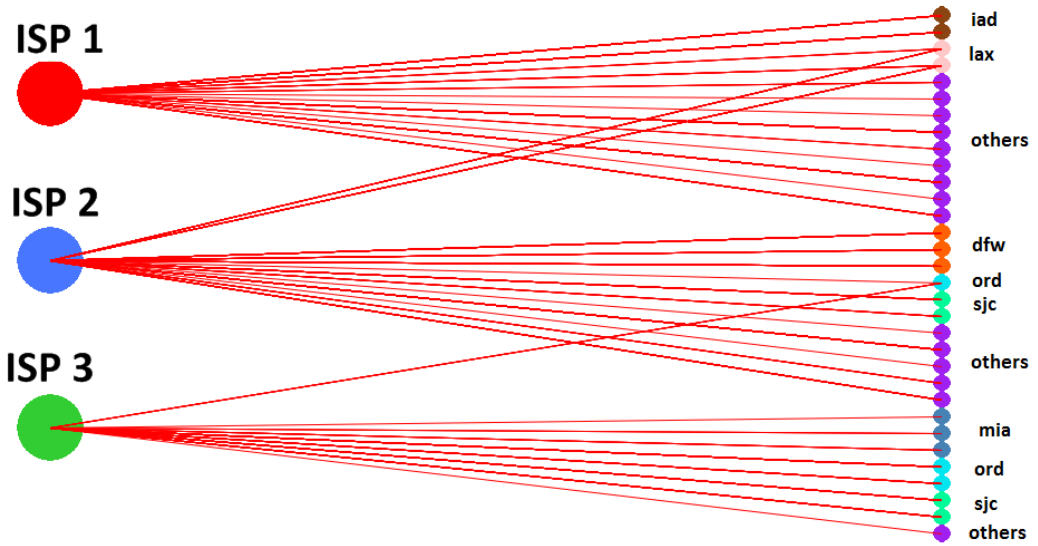

Summary

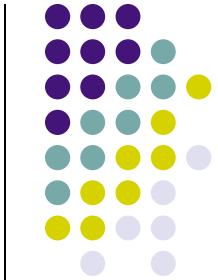

- For ISPs, Pytomo is useful to understand QoE at user level as well as server location level

- We show how to collect data to understand impact on different QoE metrics, especially in terms of location of YouTube servers/data centers.

Paper in IFIP/IEEE International Symposium on Integrated Network Management (IM'2013), Ghent, Belgium, May 2013