

Unreeling Netflix: Understanding and Improving Multi-CDN Movie Delivery

Vijay Kumar Adhikari*, Yang Guo†, Fang Hao†, Matteo Varvello†, Volker Hilt†, Moritz Steinert† and Zhi-Li Zhang*

*University of Minnesota, †Bell-Labs/Alcatel-Lucent
INFOCOM 2012

Motivation

- 23 million subscribers in US & Canada.
- Single largest source of Internet traffic in the US consuming 29.7% of downstream traffic.
- Designing a platform with high scalability and availability is challenging.
- Architecture could be used as blueprint for other providers.

Netflix Platform Overview

- Architecture is not disclosed publicly.
- Reverse engineering Netflix's architecture.
- Rely on traffic monitoring, DNS resolution, and WHOIS.

Architecture Overview

- Interact with website and monitoring all traffic during session.
- Perform DNS resolution to extract CNAMEs for all observed domain names.
- WHOIS lookup on all IP addresses to figure out owner of server.

Architecture Overview Cont.


- Three set of servers:
 - www.netflix.com or the main website hosted through Netflix's own address space.
 - Amazon's cloud mainly used for log collection authentication and DRM.
 - CDNs used for delivery of video content, Netflix uses three CDN providers namely: **Akamai, Limelight, Level3**

Architecture Overview Cont.

- Netflix utilizes the **DASH** (Dynamic Stream over HTTP) protocol for streaming. Video is encoded in different qualities and splitted up into small chunks. Player would request highest quality based on available bandwidth.

Client - Server Interaction

- Interaction between a web browser and Netflix servers during a 5 minute interval.



Client - Server Interaction Cont.

- Web player is initially downloaded from Amazon's cloud.
- A manifest file for the requested video is grabbed which contains CDN addresses along with different qualities of the video.
- Trickplay which contains thumbnails of video for seeking playback.

Client - Server Interaction Cont.

- Audio/Video chunks: downloaded more often at the beginning of playback to fill up players buffer.
- User experience reports: periodic reports are sent back to Netflix, content of messages are not discernible.

Manifest File

- XML file containing information for requested video.
- Study effect of geolocation, client capabilities and content popularity on manifest file.
- Rely on 6 user accounts and 4 computers located in different geo locations.
- Utilize 10 proxy servers running on Planet Lab nodes in different geolocations.

Manifest File - CDN Ranking

- Manifest file contains a CDN ranking which dictates the clients preference for requesting video chunks.
- No direct relation between CDN ranking and geo location or video popularity.
- CDN ranking fixed for each user over long periods of time.


Manifest File - Audio/Video Bit Rates

- Client announces computer capabilities.
- Server sends manifest file containing URL for all bit rates that client supports.
- Bitrates vary from 100 Kbps to 3600 Kbps.
- Overall 14 different bitrates are offered.

CDN Selection Strategy

- Use dummynet to throttle bandwidth.
- Initial bandwidth set to 3900 Kbps.
- Bandwidth for first CDN decreased every minute by 100Kbps.
- Once bandwidth reaches 100Kbps next CDN is throttled in the same manner.

CDN Selection Strategy - Cont.



CDN Selection Strategy - Cont.

- Client starts downloading from lower bitrates and increases bitrate by probing for higher quality videos.
- Client does not switch to other CDNs unless a lower bitrate video is not available from current CDN.

CDN Performance Measurement


- Evaluate performance of each CDN.
- Compare performance of CDNs against each other.
- Is Netflix server assignment geo aware?

CDN Performance Measurement - Methodology

- 95 vantage points spread across US.
 - 83 planet lab nodes.
 - 12 residential sites (5 different service providers).
- URL of videos extracted from manifest file.
- Measurement performed in multiple rounds.
- Each round lasts 96 seconds.
 - 4, 24 seconds slots.
 - First 3 slots allocated to each CDN - video chunk of 1.8MB downloaded.
 - In last slot requests are performed simultaneously - video chunk size of 0.6MB.
- Measurement from 8-10am on June 8th - June 26th 2012.



CDN Performance

- Average bandwidth of all providers over measurement period.
- Bandwidth for clients range from 3 Mbps to more than 200 Mbps.
- CDF curve of all CDNs are very close to each other and no provider performs significantly better than other.





CDN Performance - Cont.

- Planetlab nodes offer higher bandwidth values - last mile is bottleneck for streaming videos!
- CDN B offers higher bandwidth for couple of users - provider hosting CDN within own network.



CDN Performance - Daily Bandwidth Variation

- Compute coefficient of variance of daily average bandwidth.
- High variation in daily bandwidth for all CDNs.
- Minimum average bandwidth sufficient for SD quality.
- Similar pattern observed for hourly snapshot.


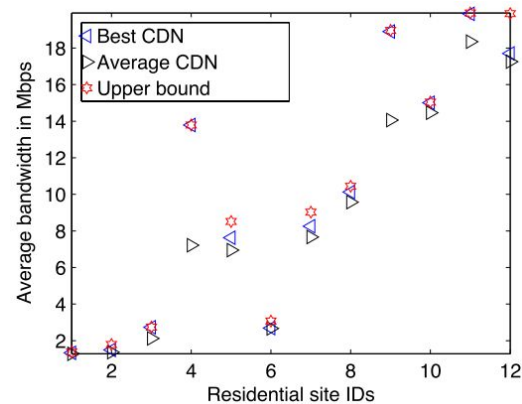


Alternate Video Delivery Strategies

- Static user assignment to CDNs could result in subpar performance for some users.
- Variation in CDN bandwidth could be utilized for better performance.
- Investigate if performance could be increased by choosing optimal CDN and if so, what strategy should be utilized.

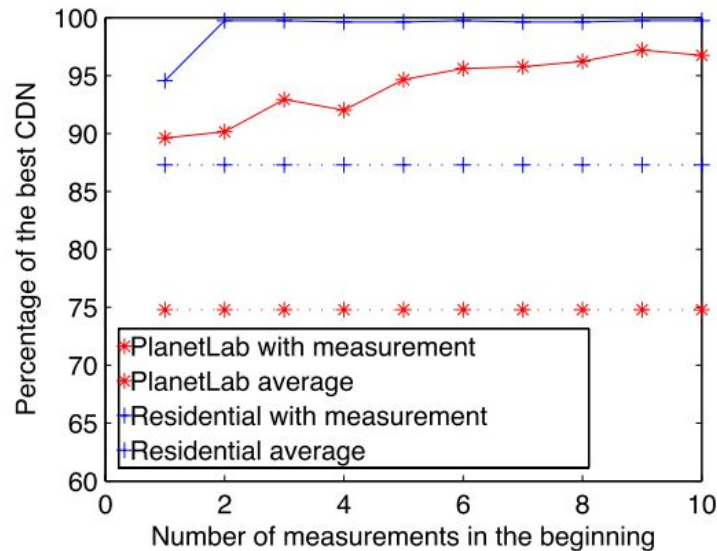
Optimal Bandwidth

- Upper bound average bandwidth reflects maximum performance given knowledge of instantaneous BW beforehand.
- Upper bound is 17%(33%) better for residential (PlanetLab) nodes.
- Best CDN offers 6%(7%) improvement.





Measurement Based CDN Selection

- Client measures BW at beginning and select top CDN for remainder of video.
- Effect of number of measurements on best CDN selection.
- 12% improvement over static assignment using more than 2 measurements.



Utilizing Multiple CDNs Simultaneously

- Measure theoretical improvement of bandwidth given the possibility of downloading chunks from multiple CDNs simultaneously.
- 54%(70%) improvement in bandwidth for residential (PlanetLab) nodes compared to top CDN.



Questions?