CSc 461/561
Multimedia Systems
Large-Scale Video Streaming

Jianping Pan

Spring 2015

Large-scale video streaming

- Over the Internet
  - short videos, long videos, etc
  - news clips, TV shows, movies, etc
  - copyrighted, user-generated, free
- Using different technologies and combinations
  - CDN or multi-CDN: YouTube, NetFlix, etc
  - P2P:e.g., PPLive, UUSee, PPStream, Joost, etc
  - P2P/cloud-assisted, hybrid, and beyond
- How to put pieces together---in the real world

Why YouTube?

World's largest (mostly user-generated) global (excl. China) video sharing service

- More than 13 million hours of video were uploaded during 2010 and 35 hours of video are uploaded every minute.
- More video is uploaded to YouTube in 60 days than the 3 major US networks created in 60 years
- 70% of YouTube traffic comes from outside the US
- YouTube reached over 700 billion playbacks in 2010
- YouTube mobile gets over 100 million views a day

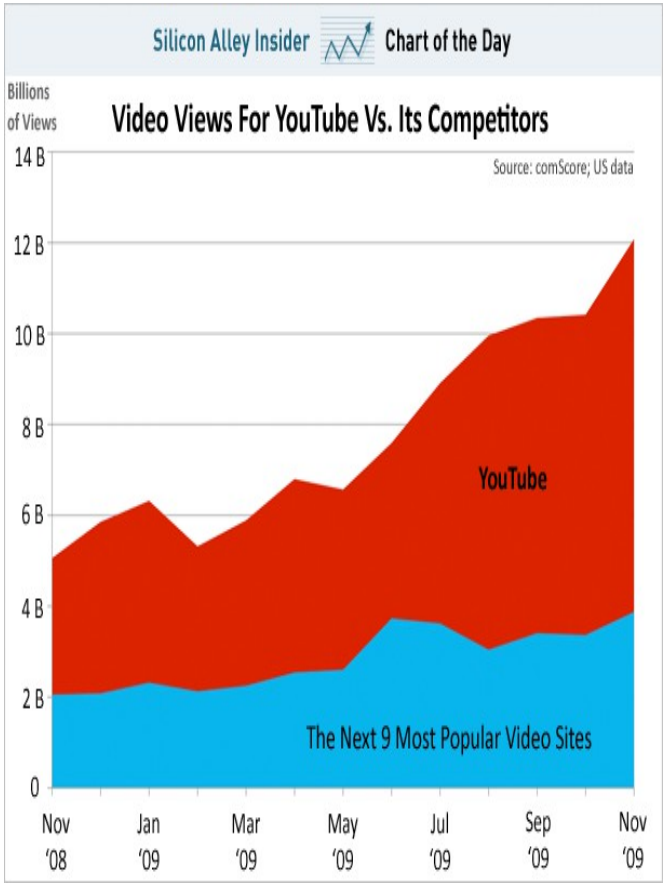

By some estimates, 5%-10% global (inter-AS) Internet traffic (2007-2009 estimate)

- up to 20% HTTP traffic (2007 estimate)

YouTube Video Delivery Basics

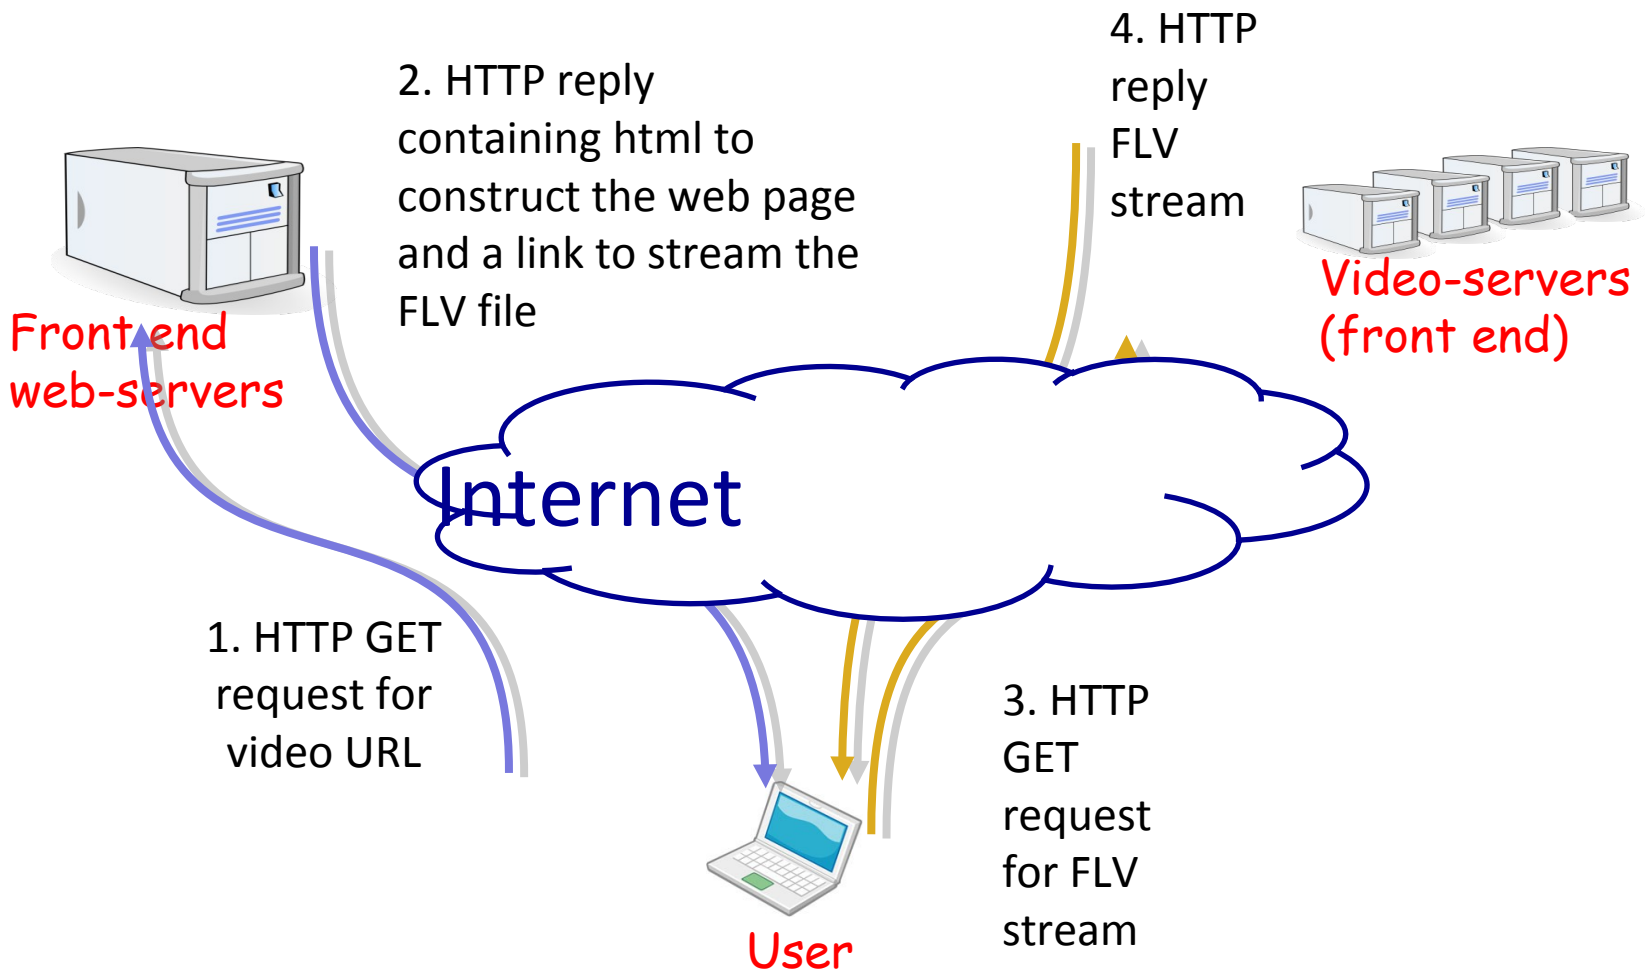

YouTube Data Center Locations
(prior Google Re-structuring)

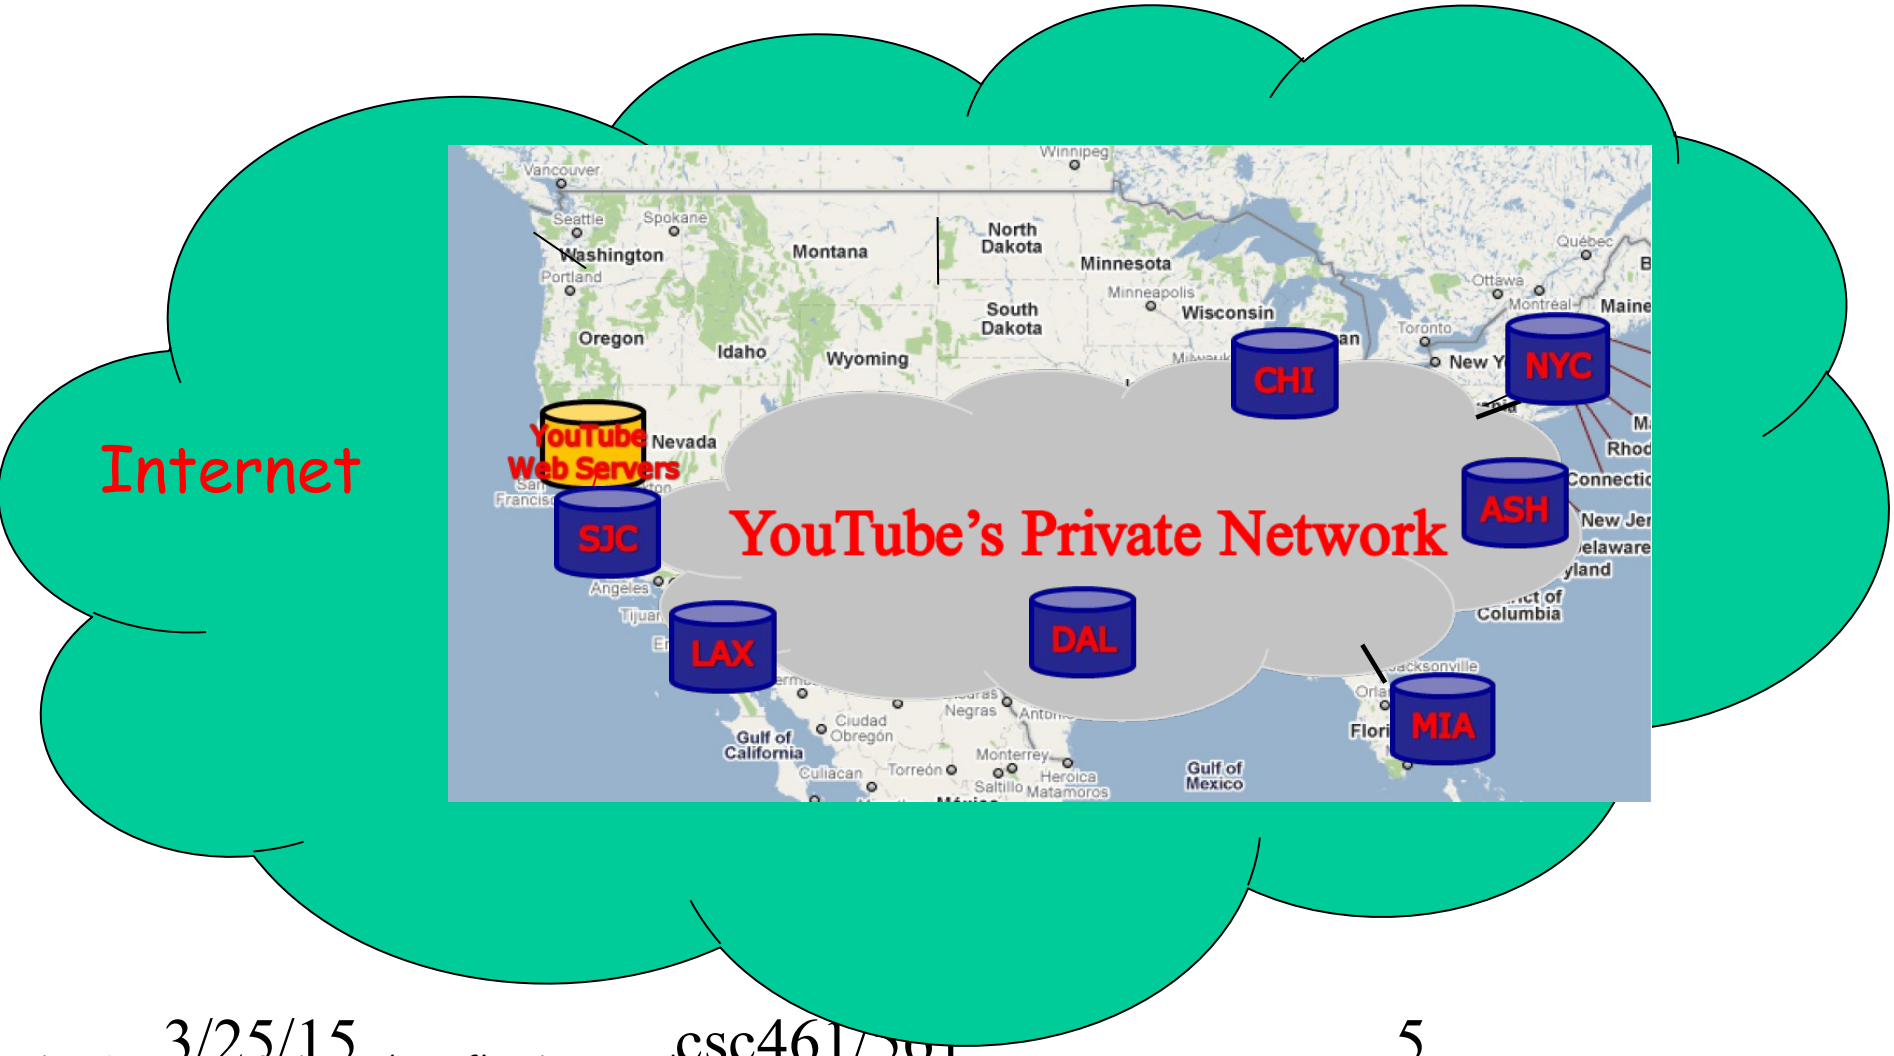

YouTube Video Id Space

- Each YouTube video is assigned a unique id
  e.g., http://www.youtube.com/watch?v=tObjCw_WgKs
- Each video id is 11 char string
  - first 10 chars can be any alpha-numeric values [0-9, a-z, A-Z] plus "-" and "_"
  - last char can be one of the 16 chars {0, 4, 8, ..., A, E, ...}

- Video id space size: 64^{11}
- Video id's are randomly distributed in the id space

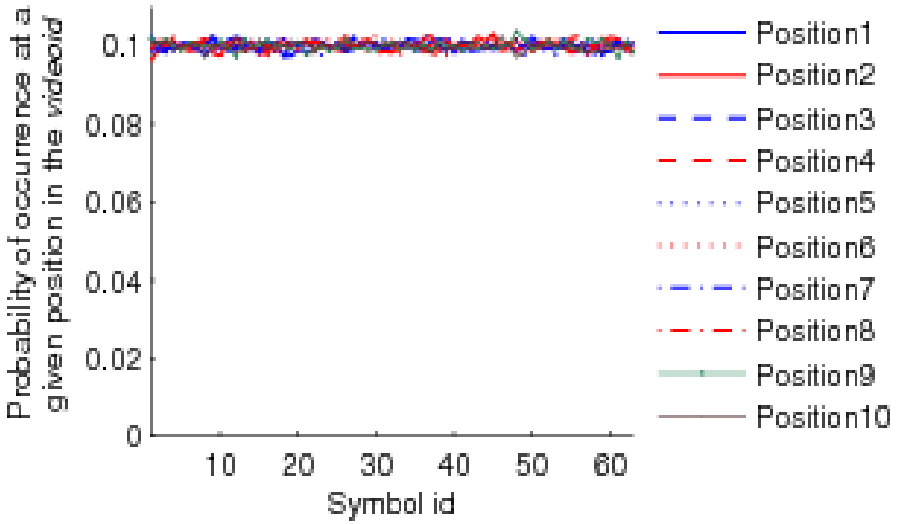

Physical Cache Hierarchy & Locations

~ 50 cache locations

- ~40 primary locations
  - including ~10 non-Google ISP locations
- 8 secondary locations
- 5 tertiary locations

Geo-locations using

- city codes in unicast hostnames, e.g., r1.**sjc**01g01.c.youtube.com
- low latency from PLnodes (< 3ms)
- clustering of IP addresses using latency matrix

YouTube Study Summary

YouTube: largest global video sharing site

"Reverse-Engineering" YouTube Delivery Cloud

comparative study of pre- vs. post-Google restructuring

Google's YouTube design provides an interesting case study of large-scale content delivery system

employs a combination of various "tricks" and mechanisms to scale with YouTube size & handle video delivery dynamics

represents some "best practice" design principles?

•[Video quality adaptation: users have to select manually!]

Interplay with ISPs and socio-technical interplay

Lessons for "future" content-centric network design

shed light on limitations of today's Internet architecture

* now youtube is DASH-like too

What Makes Netflix Interesting?

- Commercial, feature-length movies and TV shows
  - and not free; subscription-based
- Nonetheless, Netflix is huge!
  - ~25 million subscribers
  - ~20,000 titles (and growing)
  - consumes 30% of peak-time downstream bandwidth in North America
- Netflix has an interesting (cloud-sourced) architecture

	Upstream Traffic		Downstream Traffic		Total Traffic	
Rank	Application	Share	Application	Share	Application	Share
1	BitTorrent	52.01%	Netflix	29.70%	Netflix	24.71%
2	HTTP	8.31%	HTTP	18.36%	BitTorrent	17.23%
3	Skype	3.81%	YouTube	11.04%	HTTP	17.18%
4	Netflix	3.59%	BitTorrent	10.37%	YouTube	9.85%
5	PPStream	2.92%	Flash Video	4.88%	Flash Video	3.62%

SOURCE: SANDVINE NETWORK DEMOGRAPHICS

sandvine

Table 1 - North America - Top Applications by Bytes (Peak Period, Fixed Access)

What Makes Netflix Interesting?

- Commercial, feature-length movies and TV shows
  - and not free; subscription-based
- Nonetheless, Netflix is huge!
  - 25 million subscribers and ~20,000 titles (and growing)
  - consumes 30% of peak-time downstream bandwidth in North America
- A prime example of cloud-sourced architecture
  - Maintains only a small "in-house" facility for key functions
    - e.g., subscriber management (account creation, payment, …)
  - Majority of functions are sourced to Amazon cloud (EC2/S3)
    - user authentication, video search, video storage, …
  - DNS service is sourced to UltraDNS
  - Leverage multiple CDNs for video delivery
    - Akamai, Level 3 and Limelight
- Can serve as a possible blue-print for future system design
  - (nearly) "infrastructure-less" content delivery -- from Netflix's POV
  - minimize capex/opex of infrastructure, but may lose some "control" in terms of system performance, …

Netflix Architecture

- Netflix has its own "data center" for certain crucial operations (e.g., user registration, billing, …)
- Most web-based user-video interaction, computation/storage operations are cloud-sourced to Amazon AWS
- Video delivery is out/cloud-sourced to 3 CDNs
- Users need to use MS Silverlight player for video streaming

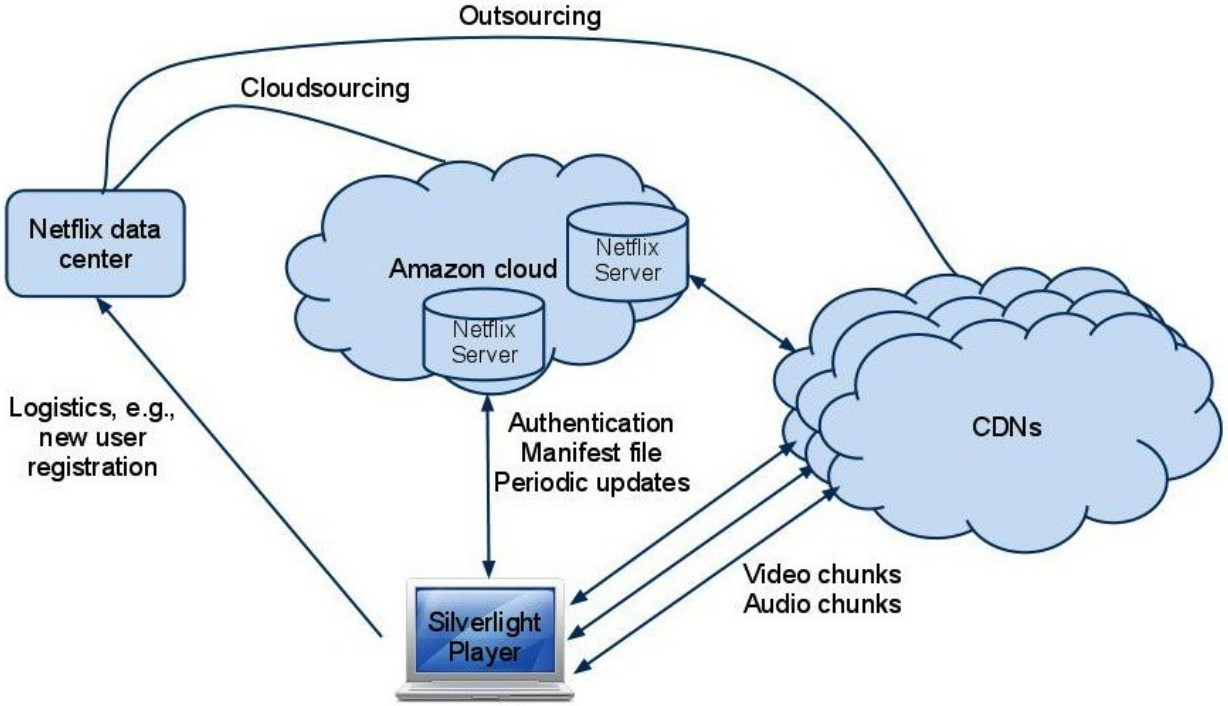

Netflix Videos and Video Chunks

- Netflix uses a numeric ID to identify each movie
  - IDs are variable length (6-8 digits): 213530, 1001192, 70221086
  - video IDs do not seem to be evenly distributed in the ID space
  - these video IDs are *not* used in playback operations

- Each movie is encoded in multiple quality levels, each is identified by a numeric ID (9 digits)
  - various numeric IDs associated with the same movie appear to have no obvious relations

Netflix Videos and CDN Namespaces

Netflix video streaming is handled directly by CDNs

- How are Netflix videos mapped to CDN namespaces & servers?

Limelight:

http://**netflix-094**.vo.llnwd.net/s/stor3/384/534975384.ismv/range/0-57689?p=58&e=1311456547&h=2caca6fb4cc2c522e657006cf69d4ace

Akamai:

http://**netflix094**.as.nflximg.com.edgesuite.net/sa53/384/534975384.ismv/range/0-57689?token=1311456547_411862e41a33dc93ee71e2e3b3fd8534

gap

Level3:

http://nflx.i.**ad483241**.x.lcdn.nflximg.com/384/534975384.ismv/range/0-57689?etime=20110723212907&movieHash=094&encoded=06847414df0656e697cbd

DASH: dynamic adaptive streaming over HTTP

Not really a protocol; it provides formats to enable efficient and high-quality delivery of streaming services over the Internet

Enable HTTP-CDNs; reuse of existing technology (codec, DRM,…)

Move "intelligence" to client: device capability, bandwidth adaptation, …

In particular, it specifies Media Presentation Description (MPD)

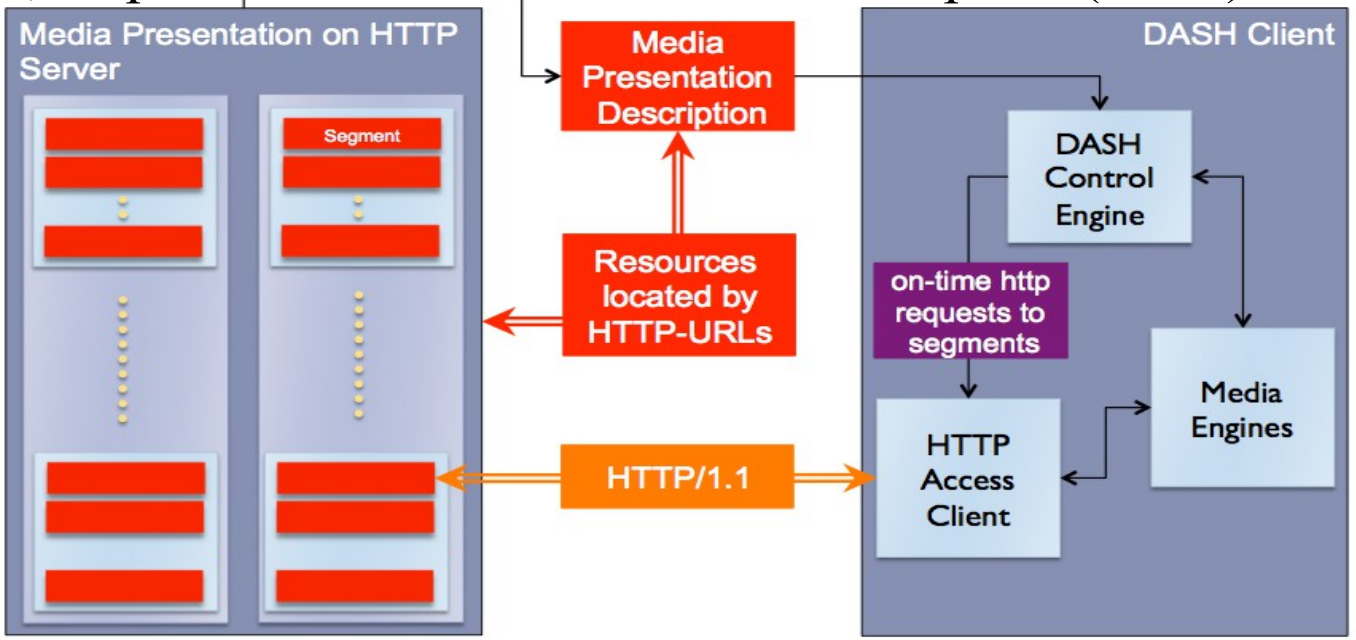

Ack & ©: Thomas Stockhammer

There is no Single Best CDN

Netflix Study Summary

- Netflix employs an interesting cloud-sourced architecture
  - Amazon AWS cloud + 3 CDNs for video streaming
- Netflix video streaming utilizes DASH
  - enables it to leverage CDNs
  - performs adaptive streaming for feature-length movies
  - allows DRM management (handled by Netflix + MS Silverlight)
- Load-balancing, cache misses or other video delivery dynamics are handled internally by each CDN
- Netflix uses a static ranking of CDNs (per user)
  - multiple CDNs are used mostly for the fail-over purpose
  - how the static ranking is determined is still a mystery (to us!)
- More "intelligent" CDN selection is possible to further improve user-perceived quality (and make them happier!)
  - As higher resolution/3D movies come into picture, these improvements will have a considerable impact