Content delivery techniques

Suzan Bayhan

Public demonstration lecture@ University of Helsinki



Target audience: Master/PhD-students in networking 3 ECTS credits

Helsinki, Finland, Feb.1st, 2017

What will you get from this course?

- How content is delivered in the current Internet? Part I
- How can it be delivered in the Future Internet? Part II
- Wireless video delivery. Part III

What will you get from this course?

- How content is delivered in the current Internet? Part I
- How can it be delivered in the Future Internet? Part II
- Wireless video delivery. Part III



- Nygren, Erik, Ramesh K. Sitaraman, and Jennifer Sun. "The Akamai network: a platform for high-performance internet applications." ACM SIGOPS Operating Systems Review 44.3 (2010): 2-19.
- Maggs, Bruce M., and Ramesh K. Sitaraman. "Algorithmic nuggets in content delivery." ACM SIGCOMM Computer Communication Review 45.3 (2015): 52-66.

Content delivery components



Content provider (Origin server)



4

Content delivery components





- Static content
 - updated infrequently
- Dynamic content
 - generated on the fly

Content delivery components









Unicast delivery not efficient large-scale delivery



Unicast delivery not efficient large-scale delivery



Not scalable for large scale content delivery!

Content delivery networks (CDN)



CDN: geographically and *strategically* distributed web servers, replicate the contents and store at the edge servers in multiple locations

Content delivery networks (CDN)



CDN: geographically distributed web servers, replicate the contents and store at the edge(replica) servers in multiple locations

CDN providers

- Emerged late 90s
- Akamai, Limelight, Level 3, MaxCDN, Incapsula, Amazon, etc.
- Akamai (1999)
 - Carrying 15-30% of web traffic
 - 216000 edge servers, 120 countries, 1500 ISPs*

CDN benefits

- Content providers (CP): delegate namespace via CNAME field in DNS or url of cdn, e.g., cdn.x.com
 - Reduced latency in distributing content
 - Lower load on CP infrastructure
 - Higher service/content availability due to multiple servers
 - Increased robustness against attacks

CDN benefits

- Content providers (CP): delegate namespace via CNAME field in DNS or url of cdn, e.g., cdn.x.com
 - Reduced latency in distributing content
 - Lower load on CP infrastructure
 - Higher service/content availability due to multiple servers
 - Increased robustness against attacks
- Users
 - Reduced latency in accessing content
 - Transparent

CDN benefits

- Content providers (CP): delegate namespace via CNAME field in DNS or url of cdn, e.g., cdn.x.com
 - Reduced latency in distributing content
 - Lower load on CP infrastructure
 - Higher service/content availability due to multiple servers
 - Increased robustness against attacks
- Users
 - Reduced latency in accessing content
 - Transparent

C Secure view-source:https://www.theguardian.com/uk-news
<source src="https://cdn.theguardian.tv/HLS/2017/01/30/170130Verhofstadt.m3u8" type="video/m3u8">
<source src="https://cdn.theguardian.tv/mainwebsite/2017/01/30/170130Verhofstadt desk.mp4" type="video/mp4">

<source src= <u>https://cdn.thequardian.tv/mainwebsite/2017/01/30/170130Verhofstadt_desk.mp4</u> type= video/mp4
<source src="https://cdn.thequardian.tv/webM/2017/01/30/170130Verhofstadt WebM.webm" type="video/webm">





• Caching (of popular content)

Popular content



- Caching
 - cache capacity is limited!
 - some content is not cacheable!



Caching

- cache capacity is limited!
- some content is not cacheable!
- Routing
 - Shorter paths
 - More reliable paths

CDN provider needs to optimize:



- Where to deploy edge servers?
 - ISP business strategies
 - Understanding of traffic demand
- Which content to replicate?
 - CP preferences
 - User preferences
- How to manage cache space?

Request routing

- Take Akamai as an example
- Which of **216000** servers in 1500 ISPs should a user get service from?



Which server is the best server for a user?



Which server is the best server for a user?

- network proximity,
- client perceived latency,
- server load
- server capacity
- content should be highly likely in the cache



Which server is the best server for a user?

- network proximity,
- client perceived latency,
- server load
- server capacity
- content should be highly likely in the cache



Akamai: User group to server cluster assignment using the stable marriage problem [Gale and Shapley, 1962]





CDN mapping system



* adapted from Chen et.al, End-User Mapping: Next Generation Request Routing for Content Delivery, ACM SIGCOMM 2015.

Shortcomings of DNS-based routing

Shortcomings of DNS-based routing

• Local DNS may not always reflect the real user location



• Local DNS may not always reflect the real user location



Figures from Chen et.al, End-User Mapping: Next Generation Request Routing for Content Delivery, ACM SIGCOMM 2015.





Local DNS may not always reflect the real user location



Figures from Chen et.al, End-User Mapping: Next Generation Request Routing for Content Delivery, ACM SIGCOMM 2015.





• Local DNS may not always reflect the real user location



Figures from Chen et.al, End-User Mapping: Next Generation Request Routing for Content Delivery, ACM SIGCOMM 2015.

(7500 Jiles)

ature: extended DNS, etc.





- Response time
- Delivery quality (for video)





- Response time
- Delivery quality (for video)



- Response time
- Delivery quality



- Response time
- Delivery quality (for video)



- User engagement
- Origin server load
- Response time
- Delivery quality



- Cache hit ratio
- Edge server utilization



- Response time
- Delivery quality (for video)



- User engagement
- Origin server load
- Response time
- Delivery quality



- Cache hit ratio
- Edge server utilization
 - Net profit





References

- CDN reference materials:
 - Nygren, Erik, Ramesh K. Sitaraman, and Jennifer Sun. "The Akamai network: a platform for high-performance internet applications." ACM SIGOPS Operating Systems Review 44.3 (2010): 2-19.
 - Dilley, John, et al. "Globally distributed content delivery." IEEE Internet Computing (2002): 50-58.
 - Maggs, Bruce M., and Ramesh K. Sitaraman. "Algorithmic nuggets in content delivery." ACM SIGCOMM Computer Communication Review 45.3 (2015): 52-66.
 - Online book: http://www.e-reading.club/bookreader.php/142069/Buyya_-_Content_Delivery_Networks.pdf
- ICN reference materials:
 - Ahlgren, Bengt, et al. "A survey of information-centric networking." IEEE Communications Magazine 50.7 (2012).
 - Xylomenos, George, et al. "A survey of information-centric networking research." IEEE Communications Surveys & Tutorials 16.2 (2014): 1024-1049.
- Video delivery reading list:
 - Thomas Stockhammer. 2011. Dynamic adaptive streaming over HTTP: standards and design principles. ACM MMSys.
 - Begen et al. "Watching video over the Web, part 2: applications, standardization, and open issues," IEEE Internet Computing, May/June 2011
 - Begen et al., "Watching video over the Web, part 1: streaming protocols," IEEE Internet Computing, Mar./Apr. 2011