

CS640: Introduction to Computer Networks

Aditya Akella

Lecture 18 -
The Web, Caching and CDNs

Announcements

- PA 3 out today
- HW 4 will be out some time next week
- HW 3 due today
- Sign up for PA 2 demo slot soon!

2

The Road Ahead

- HTTP and TCP
- HTTP caching
- Content distribution networks

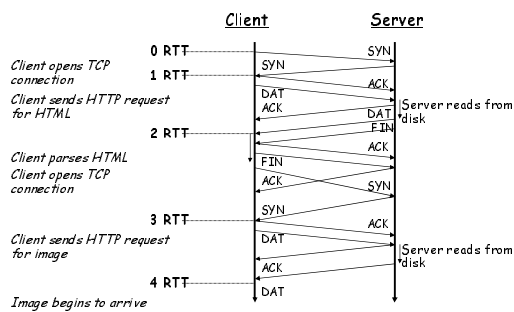
3

HTTP 0.9/1.0

- One request/response per TCP connection
 - Simple to implement
- Disadvantages
 - Multiple connection setups → three-way handshake each time
 - Several extra round trips added to transfer
 - Multiple slow starts
 - Why is this bad?

4

Single Transfer Example



5

More Problems

- Short transfers are hard on TCP
 - Stuck in slow start
 - Also, loss recovery is poor when windows are small
- Lots of extra connections
 - Increases server state/processing
- Server also forced to keep TIME_WAIT connection state
 - Tends to be an order of magnitude greater than # of active connections

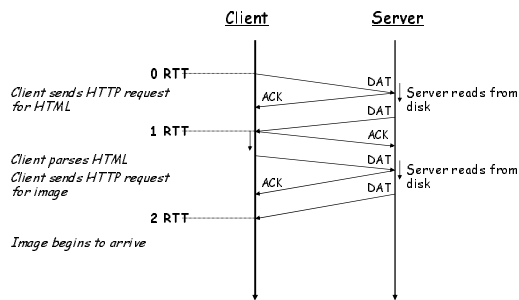
6

Persistent Connection Solution

- Multiplex multiple transfers onto one TCP connection
- How to identify requests/responses
 - Delimiter → Server must examine response for delimiter string
 - Content-length and delimiter → Must know size of transfer in advance
 - Block-based transmission → send in multiple length-delimited blocks
 - Store-and-forward → wait for entire response and then use content-length
 - Solution → use existing methods and close connection otherwise

7

Persistent Connection Example



8

Persistent HTTP

Nonpersistent HTTP issues:

- Requires 2 RTTs per object
- OS must work and allocate host resources for each TCP connection
- But browsers often open parallel TCP connections to fetch referenced objects

Persistent HTTP

- Server leaves connection open after sending response
- Subsequent HTTP messages between same client/server are sent over connection

Persistent without pipelining:

- Client issues new request only when previous response has been received
- One RTT for each referenced object

Persistent with pipelining:

- Default in HTTP/1.1
- Client sends requests as soon as it encounters a referenced object
- As little as one RTT for all the referenced objects

9

HTTP Caching

- Why caching?
- Clients often cache documents
 - Challenge: update of documents
 - If-Modified-Since requests to check
 - HTTP 0.9/1.0 used just date
 - HTTP 1.1 has an opaque "entity tag" (could be a file signature, etc.) as well
- When/how often should the original be checked for changes?
 - Check every time?
 - Check each session? Day? Etc?
 - Use "Expires" header
 - If no Expires, often use Last-Modified as estimate

10

Example Cache Check Request

```
GET / HTTP/1.1
Accept: */*
Accept-Language: en-us
Accept-Encoding: gzip, deflate
If-Modified-Since: Mon, 29 Jan 2001 17:54:18
GMT
If-None-Match: "7a11f-10ed-3a75ae4a"
User-Agent: Mozilla/4.0 (compatible; MSIE 5.5;
Windows NT 5.0)
Host: www.intel-iris.net
Connection: Keep-Alive
```

11

Example Cache Check Response

```
HTTP/1.1 304 Not Modified
Date: Tue, 27 Mar 2001 03:50:51 GMT
Server: Apache/1.3.14 (Unix) (Red-Hat/Linux)
mod_ssl/2.7.1 OpenSSL/0.9.5a DAV/1.0.2
PHP/4.0.1pl2 mod_perl/1.24
Connection: Keep-Alive
Keep-Alive: timeout=15, max=100
ETag: "7a11f-10ed-3a75ae4a"
```

12


Caching Example (1)

Assumptions

- Average object size = 100,000 bits
- Avg. request rate from institution's browser to origin servers = 15/sec
- Delay from institutional router to any origin server and back to router = 2 sec

Consequences

- Utilization on LAN = 15%
- Utilization on access link = 100%
- Total delay = Internet delay + access delay + LAN delay = 2 sec + minutes + milliseconds



13


Caching Example (2)

Possible solution

- Increase bandwidth of access link to, say, 10 Mbps
- Often a costly upgrade

Consequences

- Utilization on LAN = 15%
- Utilization on access link = 15%
- Total delay = Internet delay + access delay + LAN delay = 2 sec + msec + msec



14


Caching Example (3)

Install cache

- Suppose hit rate is .4

Consequence

- 40% requests will be satisfied almost immediately (say 10 msec)
- 60% requests satisfied by origin server
- Utilization of access link reduced to 60%, resulting in negligible delays
- Weighted average of delays = $.6 * 2 \text{ sec} + .4 * 10 \text{ msec} < 1.3 \text{ secs}$




15

Web Proxy Caches

- User configures browser: Web accesses via cache

- Browser sends all HTTP requests to cache

- Object in cache: cache returns object
- Else cache requests object from origin server, then returns object to client



16

Problems

- Over 50% of all HTTP objects are uncacheable - why?

- Not easily solvable

- Dynamic data → stock prices, scores, web cams
- CGI scripts → results based on passed parameters
- SSL → encrypted data is not cacheable
 - Most web clients don't handle mixed pages well → many generic objects transferred with SSL
- Cookies → results may be based on passed data
- Hit metering → owner wants to measure # of hits for revenue, etc.

- What will be the end result?

17

Content Distribution Networks & Server Selection

- Replicate content on many servers

- Challenges

- Which content to replicate
- How to replicate content
- Where to place replicas
- How to find replicated content
- How to choose among known replicas
- How to direct clients towards replica

18

Server Selection

- Which server?
 - Lowest load → to balance load on servers
 - Best performance → to improve client performance
 - Based on Geography? RTT? Throughput? Load?
 - Any alive node → to provide fault tolerance
- How to direct clients to a particular server?
 - As part of routing → anycast, cluster load balancing
 - Not covered today...
 - As part of application → HTTP redirect
 - As part of naming → DNS

19

Application-Based Redirection

- HTTP supports simple way to indicate that Web page has moved (30X responses)
- Server receives Get request from client
 - Decides which server is best suited for particular client and object
 - Returns HTTP redirect to that server
- Can make informed application specific decision
- May introduce additional overhead → multiple connection setup, name lookups, etc.

20

Naming Based

- Client does name lookup for service
- Name server chooses appropriate server address
 - A-record returned is "best" one for the client
- What information can name server base decision on?
 - Server load/location → must be collected
 - Information in the name lookup request
 - Name service client → typically the local name server for client


21

Content Distribution Networks (CDNs)

- The content providers are the CDN customers.

Content replication

- CDN company installs hundreds of CDN servers throughout Internet
 - Close to users
- CDN replicates its customers' content in CDN servers. When provider updates content, CDN updates servers



22

How Akamai Works

- Clients fetch html document from primary server
 - E.g. fetch index.html from cnn.com
- "Akamaized" URLs for replicated content are replaced in html
 - E.g. `` replaced with ``
- Client is forced to resolve `aXYZ.g.akamai.net` hostname

23

How Akamai Works

- How is content replicated?
 - Akamai only replicates static content (*)
 - Modified name contains original file name and content provider ID
 - Akamai server is asked for content
 - First checks local cache
 - If not in cache, requests file from primary server; caches file
- * (At least, the version we're talking about today. Akamai actually lets sites write code that can run on Akamai's servers, but that's a different beast altogether!)

24

How Akamai Works

- Root server gives NS record for akamai.net
- Akamai.net name server returns NS record for g.akamaitech.net
 - Name server chosen to be in region of client's name server
 - Out-of-band measurements to obtain this
- G.akamaitech.net nameserver chooses server in region
 - Which server to choose?
 - Uses aXYZ name and hash

25

Simple Hashing

- Given document XYZ, we need to choose a server to use
- Suppose we use modulo
- Number servers from 1...n
 - Place document XYZ on server (XYZ mod n)
 - What happens when a servers fails? $n \rightarrow n-1$
 - Why might this be bad?

26

Consistent Hash


- Desired features
 - Balanced - load is equal across buckets
 - Smoothness - little impact on hash bucket contents when buckets are added/removed
 - Spread - small set of hash buckets that may hold a set of object
 - Load - # of objects assigned to hash bucket is small

27

Consistent Hash - Example

- Construction

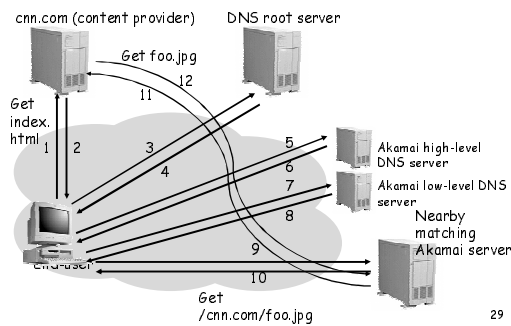
- Assign each of C hash buckets to random points on mod 2^n circle, where, hash key size = n .
- Map object to random position on circle
- Hash of object = closest clockwise bucket



- Smoothness → addition of bucket does not cause movement between existing buckets
- Spread & Load → small set of buckets that lie near object
- Balance → no bucket is responsible for large number of objects


28

How Akamai Works



29

Akamai - Subsequent Requests



30

Summary

- **HTTP: Simple text-based file exchange protocol**
 - Support for status/error responses, authentication, client-side state maintenance, cache maintenance
- **Workloads**
 - Typical documents structure, popularity
 - Server workload
- **Interactions with TCP**
 - Connection setup, reliability, state maintenance
 - Persistent connections
- **How to improve performance**
 - Persistent connections
 - Caching
 - Replication

31
