



15-441 Computer Networking

Lecture 6 – Web Optimizations

Outline



- Persistent HTTP
- HTTP Caching
- Server Selection & Content Distribution Networks

Typical Workload (Web Pages)



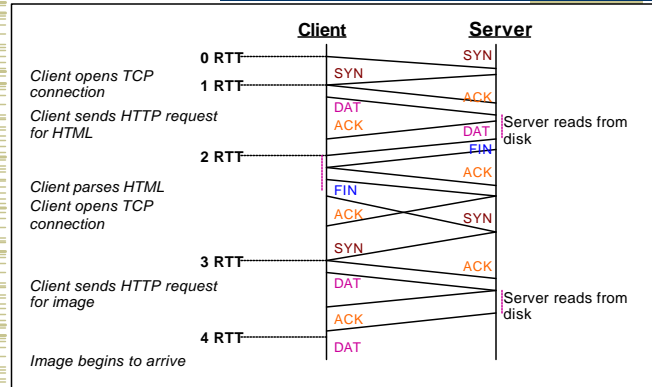
- Multiple (typically small) objects per page
- File sizes
 - Heavy-tailed
 - Pareto distribution for tail
 - Lognormal for body of distribution
- Embedded references
 - Number of embedded objects =
pareto – $p(x) = ak^ax^{(a+1)}$

HTTP 0.9/1.0



- One request/response per TCP connection
 - Simple to implement
 - Uses connection close to delimit objects
- Disadvantages
 - Multiple connection setups → three-way handshake each time
 - Several extra round trips added to transfer
 - Multiple slow starts

Single Transfer Example



Lecture 6: 09-16-2002

5

More Problems



- Short transfers are hard on TCP
 - Stuck in slow start
 - 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
 - Why must server keep these?
 - Tends to be an order of magnitude greater than # of active connections, why?

Lecture 6: 09-16-2002

6

Netscape Solution



- Mosaic (original popular Web browser) fetched one object at a time!
- Netscape uses multiple concurrent connections to improve response time
 - Different parts of Web page arrive independently
 - Can grab more of the network bandwidth than other users
- Doesn't necessarily improve response time
 - TCP loss recovery ends up being timeout dominated because windows are small

Lecture 6: 09-16-2002

7

Persistent Connection Solution

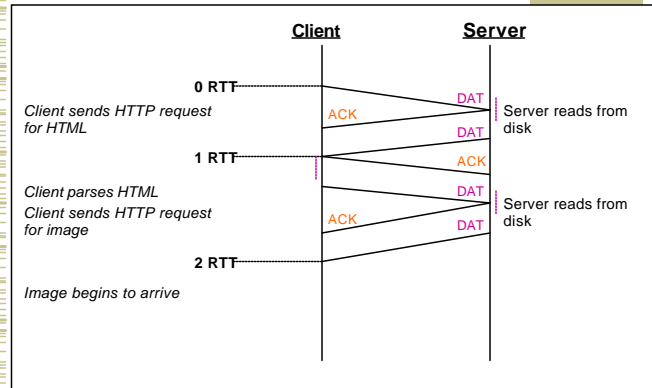


- Multiplex multiple transfers onto one TCP connection
- How to identify requests/responses
 - Delimiter → Server must examine response for delimiter string
 - Content-length and delimiter → Must know size of transfer in advance
 - Block-based transmission → send in multiple length delimited blocks
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

Lecture 6: 09-16-2002

8

Persistent Connection Example



Lecture 6: 09-16-2002

9

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

Lecture 6: 09-16-2002

10

Persistent Connection Performance



- Benefits greatest for small objects
 - Up to 2x improvement in response time
- Server resource utilization reduced due to fewer connection establishments and fewer active connections
- TCP behavior improved
 - Longer connections help adaptation to available bandwidth
 - Larger congestion window improves loss recovery

Lecture 6: 09-16-2002

11

Remaining Problems



- Serialized transmission
 - Stall in transfer of one object prevents delivery of others
 - Much of the useful information in first few bytes
 - Can "packetize" transfer over TCP
 - Could use range requests
- Application specific solution to transport protocol problems
 - Solve the problem at the transport layer
 - Could fix TCP so it works well with multiple simultaneous connections
 - More difficult to deploy

Lecture 6: 09-16-2002

12

Outline



- Persistent HTTP
- **HTTP Caching**
- Server Selection & Content Distribution Networks

Lecture 6: 09-16-2002

13

Typical Workload (Server)



- Popularity
 - Zipf distribution ($P = kr^{-1}$) → surprisingly common
 - Obvious optimization → caching
- Request sizes
 - In one measurement paper → median 1946 bytes, mean 13767 bytes
 - Why such a difference? Heavy-tailed distribution
 - Pareto – $p(x) = ak^{ax(a+1)}$
- Temporal locality
 - Modeled as distance into push-down stack
 - Lognormal distribution of stack distances
- Request interarrival
 - Bursty request patterns

Lecture 6: 09-16-2002

14

HTTP 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 file signature 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

Lecture 6: 09-16-2002

15

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
```

Lecture 6: 09-16-2002

16

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"


Lecture 6: 09-16-2002

17

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



Lecture 6: 09-16-2002

18

Proxy Caching



- Goal: Satisfy client request without involving origin server
 - Reduce client response time
 - Reduce network bandwidth usage
 - Wide area vs. local area use
 - These two objectives are often in conflict
 - May do exhaustive local search to avoid using wide area bandwidth
 - Prefetching uses extra bandwidth to reduce client response time

Lecture 6: 09-16-2002

19

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



Lecture 6: 09-16-2002

20

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



Lecture 6: 09-16-2002

21

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
= $.6 * 2 \text{ sec} + .4 * 10 \text{ msec} < 1.3 \text{ secs}$



Lecture 6: 09-16-2002

22

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
- Obvious fixes
 - 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?

Lecture 6: 09-16-2002

23

Caching Proxies – Sources for Misses



- Capacity
 - How large a cache is necessary or equivalent to infinite
 - On disk vs. in memory → typically on disk
- Compulsory
 - First time access to document
 - Non-cacheable documents
 - CGI-scripts
 - Personalized documents (cookies, etc)
 - Encrypted data (SSL)
- Consistency
 - Document has been updated/expired before reuse
- Conflict
 - No such misses

Lecture 6: 09-16-2002

24

Proxy Implementation Problems



- Aborted transfers
 - Many proxies transfer entire document even though client has stopped → eliminates saving of bandwidth
- Making objects cacheable
 - Proxy's apply heuristics → cookies don't apply to some objects, guesswork on expiration
 - May not match client behavior/desires
- Client misconfiguration
 - Many clients have either absurdly small caches or no cache
- How much would hit rate drop if clients did the same things as proxies

Outline

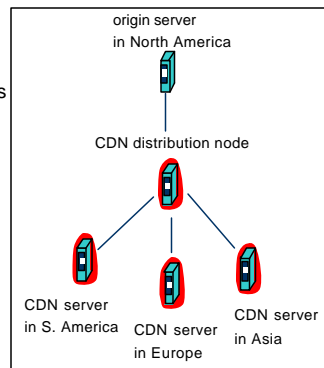


- Persistent HTTP
- HTTP Caching
- **Server Selection & Content Distribution Networks**

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



Content Distribution Networks & Server Selection



- Replicate content on many servers
- Challenges
 - How to replicate content
 - Where to replicate content
 - How to find replicated content
 - How to choose among known replicas
 - How to direct clients towards replica

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 ☹
 - As part of application → HTTP redirect
 - As part of naming → DNS

Lecture 6: 09-16-2002

29

Application Based



- 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.
- While good solution in general, but...
 - HTTP Redirect has some design flaws – especially with current browsers

Lecture 6: 09-16-2002

30

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

Lecture 6: 09-16-2002

31

Naming Based



- Round-robin
 - Randomly choose replica
 - Avoid hot-spots
- [Semi-]static metrics
 - Geography
 - Route metrics
 - How well would these work?
- Predicted application performance
 - How to predict?
 - Only have limited info at name resolution

Lecture 6: 09-16-2002

32

How Akamai Works



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

Lecture 6: 09-16-2002

33

How Akamai Works



- How is content replicated?
- Akamai only replicates static content
- Modified name contains original file name
- Akamai server is asked for content
 - First checks local cache
 - If not in cache, requests file from primary server and caches file

Lecture 6: 09-16-2002

34

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
 - TTL is large
- `G.akamaitech.net` nameserver chooses server in region
 - Should try to choose server that has file in cache - How to choose?
 - Uses `aXYZ` name and hash
 - TTL is small → why?

Lecture 6: 09-16-2002

35

Simple Hashing



- Given document `XYZ`, we need to choose a server to use
- Suppose we use modulo
- Number servers from $1 \dots n$
 - Place document `XYZ` on server $(XYZ \bmod n)$
 - What happens when a server fails? $n \rightarrow n-1$
 - Same if different people have different measures of n
 - Why might this be bad?

Lecture 6: 09-16-2002

36

Consistent Hash



- “view” = subset of all hash buckets that are visible
- Desired features
 - Balanced – in any one view, 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 an object regardless of views
 - Load – across all views # of objects assigned to hash bucket is small

Lecture 6: 09-16-2002


37

Consistent Hash – Example



Construction

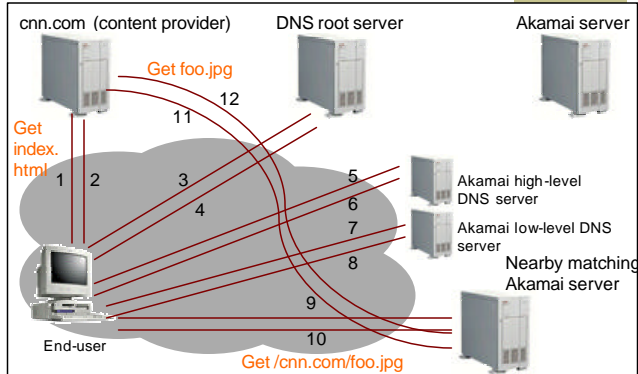
- Assign each of C hash buckets to random points on mod 2^n circle, where, hash key size = n .
- Map object to random position on circle
- Hash of object = closest clockwise bucket
- Smoothness → addition of bucket does not cause movement between existing buckets
- Spread & Load → small set of buckets that lie near object
- Balance → no bucket is responsible for large number of objects



Lecture 6: 09-16-2002

38

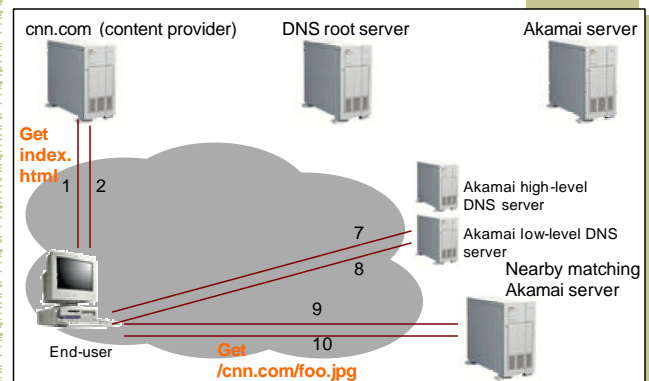
How Akamai Works



Lecture 6: 09-16-2002

39

Akamai – Subsequent Requests



Lecture 6: 09-16-2002

40

Impact on DNS Usage



- DNS is used for server selection more and more
 - What are reasonable DNS TTLs for this type of use
 - Typically want to adapt to load changes
 - Low TTL for A-records → what about NS records?
- How does this affect caching?
- What do the first and subsequent lookup do?

HTTP (Summary)



- 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

Next Lecture



- Transport introduction
- Error recovery
- TCP flow control
- TCP connection setup/data transfer