

Computer Network Applications

Lecture 6

Fall 2008


Dr. Hui Xiong
Rutgers University

Introduction 1-1

Last Course Review: Web caches (proxy server)

Goal: satisfy client request without involving origin server

- user sets 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



Introduction 1-2

More about Web caching

- Cache acts as both client and server
 - Typically cache is installed by ISP (university, company, residential ISP)
- Why Web caching?**
- Reduce response time for client request.
 - Reduce traffic on an institution's access link.
 - Internet dense with caches enables "poor" content providers to effectively deliver content (but so does P2P file sharing)

Introduction 1-3

Cookies

What cookies can bring:

- authorization
- shopping carts
- recommendations
- user session state (Web e-mail)

- aside**
- Cookies and privacy:**
- cookies permit sites to learn a lot about you
 - you may supply name and e-mail to sites
 - search engines use redirection & cookies to learn yet more
 - advertising companies obtain info across sites

Introduction 1-4

Chapter 2: Application layer

- 2.1 Principles of network applications
- 2.2 Web and HTTP
- 2.3 FTP
- 2.4 Electronic Mail
 - SMTP, POP3, IMAP
- 2.5 DNS
- 2.6 P2P file sharing
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP
- 2.9 Building a Web server

Introduction 1-5

DNS: Domain Name System

People: many identifiers:

- SSN, name, passport #

Internet hosts, routers:

- IP address (32 bit) - used for addressing datagrams
- "name", e.g., ww.yahoo.com - used by humans

Q: map between IP addresses and name ?

Domain Name System:

- *distributed database* implemented in hierarchy of many *name servers*
- *application-layer protocol* host, routers, name servers to communicate to *resolve* names (address/name translation)
 - note: core Internet function, implemented as application-layer protocol
 - complexity at network's "edge"

Introduction 1-6

DNS

DNS services

- ❑ Hostname to IP address translation
- ❑ Host aliasing
 - Canonical and alias names
- ❑ Mail server aliasing
- ❑ Load distribution
 - Replicated Web servers: set of IP addresses for one canonical name


Why not centralize DNS?

- ❑ single point of failure
- ❑ traffic volume
- ❑ distant centralized database
- ❑ maintenance

doesn't *scale!*

Introduction 1-7

Distributed, Hierarchical Database



Client wants IP for www.amazon.com; 1st approx:

- ❑ Client queries a root server to find com DNS server
- ❑ Client queries com DNS server to get amazon.com DNS server
- ❑ Client queries amazon.com DNS server to get IP address for www.amazon.com

Introduction 1-8

DNS: Root name servers

- ❑ contacted by local name server that can not resolve name
- ❑ root name server:
 - contacts authoritative name server if name mapping not known
 - gets mapping
 - returns mapping to local name server



13 root name servers worldwide

Introduction 1-9

TLD and Authoritative Servers

- ❑ **Top-level domain (TLD) servers:** responsible for com, org, net, edu, etc, and all top-level country domains uk, fr, ca, cn.
 - Network solutions maintains servers for com TLD
 - Educause for edu TLD
- ❑ **Authoritative DNS servers:** organization's DNS servers, providing authoritative hostname to IP mappings for organization's servers (e.g., Web and mail).
 - Can be maintained by organization or service provider

Introduction 1-10


Local Name Server

- ❑ Does not strictly belong to hierarchy
- ❑ Each ISP (residential ISP, company, university) has one.
 - Also called "default name server"
- ❑ When a host makes a DNS query, query is sent to its local DNS server
 - Acts as a proxy, forwards query into hierarchy.

Introduction 1-11

Example

- ❑ Host at cis.poly.edu wants IP address for gaia.cs.umass.edu



Introduction 1-12


Recursive queries

recursive query:

- puts burden of name resolution on contacted name server
- heavy load?

iterated query:

- contacted server replies with name of server to contact
- "I don't know this name, but ask this server"



Introduction 1-13

DNS: caching and updating records

- once (any) name server learns mapping, it *caches* mapping
 - cache entries timeout (disappear) after some time
 - TLD servers typically cached in local name servers
 - Thus root name servers not often visited
- update/notify mechanisms under design by IETF
 - RFC 2136
 - <http://www.ietf.org/html.charters/dnsind-charter.html>

Introduction 1-14

DNS records

DNS: distributed db storing resource records (RR)

RR format: (name, value, type, ttl)

- Type=A
 - name is hostname
 - value is IP address
- Type=NS
 - name is domain (e.g. foo.com)
 - value is IP address of authoritative name server for this domain
- Type=CNAME
 - name is alias name for some "canonical" (the real) name
 - www.ibm.com is really servereast.backup2.ibm.com
 - value is canonical name
- Type=MX
 - value is name of mailserver associated with name

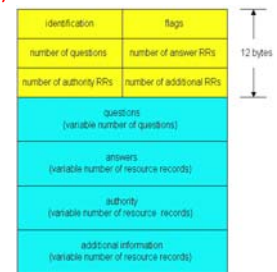
Introduction 1-15

DNS protocol, messages

DNS protocol: *query* and *reply* messages, both with same *message format*


msg header

- identification: 16 bit # for query, reply to query uses same #
- flags:
 - query or reply
 - recursion desired
 - recursion available
 - reply is authoritative



Introduction 1-16

DNS protocol, messages



Introduction 1-17

Inserting records into DNS

- Example: just created startup "Network Utopia"
- Register name networkkuptopia.com at a **registrar** (e.g., Network Solutions)
 - Need to provide registrar with names and IP addresses of your authoritative name server (primary and secondary)
 - Registrar inserts two RRs into the com TLD server:


```
(networkkuptopia.com, dns1.networkkuptopia.com, NS)
(dns1.networkkuptopia.com, 212.212.212.1, A)
```
- Put in authoritative server Type A record for www.networkkuptopia.com and Type MX record for networkkuptopia.com
- How do people get the IP address of your Web site?

Introduction 1-18

Chapter 2: Application layer

- 2.1 Principles of network applications
 - app architectures
 - app requirements
- 2.2 Web and HTTP
- 2.4 Electronic Mail
 - SMTP, POP3, IMAP
- 2.5 DNS
- 2.6 P2P file sharing

Introduction 1-19

P2P file sharing

Example


- Alice runs P2P client application on her notebook computer
 - Intermittently connects to Internet; gets new IP address for each connection
 - Asks for "Hey Jude"
 - Application displays other peers that have copy of Hey Jude.
 - Alice chooses one of the peers, Bob.
 - File is copied from Bob's PC to Alice's notebook: HTTP
 - While Alice downloads, other users uploading from Alice.
 - Alice's peer is both a Web client and a transient Web server.
- All peers are servers = highly scalable!

Introduction 1-20

P2P: centralized directory

original "Napster" design

- 1) when peer connects, it informs central server:
 - IP address
 - content
- 2) Alice queries for "Hey Jude"
- 3) Alice requests file from Bob



Introduction 1-21

P2P: problems with centralized directory

- Single point of failure
- Performance bottleneck
- Copyright infringement

file transfer is decentralized, but locating content is highly centralized

Introduction 1-22

Query flooding: Gnutella


- fully distributed
 - no central server
- public domain protocol
- many Gnutella clients implementing protocol
- **overlay network: graph**
 - edge between peer X and Y if there's a TCP connection
 - all active peers and edges is overlay net
 - Edge is not a physical link
 - Given peer will typically be connected with < 10 overlay neighbors

Introduction 1-23

Gnutella: protocol

- Query message sent over existing TCP connections
- peers forward Query message
- QueryHit sent over reverse path
- File transfer: HTTP

Scalability: limited scope flooding



Introduction 1-24

Gnutella: Peer joining


1. Joining peer X must find some other peer in Gnutella network: use list of candidate peers
2. X sequentially attempts to make TCP with peers on list until connection setup with Y
3. X sends Ping message to Y; Y forwards Ping message.
4. All peers receiving Ping message respond with Pong message
5. X receives many Pong messages. It can then setup additional TCP connections

Peer leaving: see homework problem!

Introduction 1-25

Exploiting heterogeneity: KaZaA

- Each peer is either a group leader or assigned to a group leader.
 - TCP connection between peer and its group leader.
 - TCP connections between some pairs of group leaders.
- Group leader tracks the content in all its children.



Introduction 1-26

KaZaA: Querying

- Each file has a hash and a descriptor
- Client sends keyword query to its group leader
- Group leader responds with matches:
 - For each match: metadata, hash, IP address
- If group leader forwards query to other group leaders, they respond with matches
- Client then selects files for downloading
 - HTTP requests using hash as identifier sent to peers holding desired file

Introduction 1-27

Kazaa tricks

- Limitations on simultaneous uploads
- Request queuing
- Incentive priorities
- Parallel downloading

Introduction 1-28

Chapter 2: Summary

Our study of network apps now complete!

- Application architectures
 - client-server
 - P2P
 - hybrid
- application service requirements:
 - reliability, bandwidth, delay
- Internet transport service model
 - connection-oriented, reliable: TCP
 - unreliable, datagrams: UDP
- specific protocols:
 - HTTP
 - FTP
 - SMTP, POP, IMAP
 - DNS

Introduction 1-29

Chapter 2: Summary

Most importantly: learned about protocols

- typical request/reply message exchange:
 - client requests info or service
 - server responds with data, status code
- message formats:
 - headers: fields giving info about data
 - data: info being communicated
- control vs. data msgs
 - in-band, out-of-band
- centralized vs. decentralized
- stateless vs. stateful
- reliable vs. unreliable msg transfer
- "complexity at network edge"

Introduction 1-30