UDP Server

import java.io.*; import java.net.*; class UDPServer {

public static void main(String argv[]) throws Exception
{

String sentence; String capitalizedSentence; byte[] receiveData = new byte[1024]; byte[] sendData = new byte[1024]; DatagramSocket serverSocket = new DatagramSocket(9876);

while (true) {

DatagramPacket receivePacket = new DatagramPacket(receiveData, receiveData.length); serverSocket.receive(receivePacket); sentence = new String(receivePacket.getData()); InetAddress IPAddress = receivePacket.getAddress(); int port = receivePacket.getPort(); capitalizedSentence = sentence.toUpperCase(); sendData = capitalizedSentence.getBytes(); DatagramPacket sendPacket = new DatagramPacket(sendData, sendData.length, IPAddress, port); serverSocket.send(sendPacket);

} } }

Application Layer 2-1



import java.io.*; import java.net.*; import java.util.Scanner; class UDPClient {

public static void main(String argv[]) throws Exception

Scanner kbd = new Scanner(System.in); String sentence; String modifiedSentence; DatagramSocket clientSocket = new DatagramSocket(); InetAddress IPAddress = InetAddress.getByName("localhost"); byte[] sendData = new byte[1024]; byte[] receiveData = new byte[1024];

System.out.println("Enter some text."); sentence = kbd.nextLine();

sendData = sentence.getBytes(); DatagramPacket sendPacket = new DatagramPacket(sendData, sendData.length,IPAddress, 9876);

clientSocket.send(sendPacket); DatagramPacket receivePacket = new DatagramPacket(receiveData, receiveData.length);

clientSocket.receive(receivePacket); modifiedSentence = new String(receivePacket.getData());

System.out.println("From Server: " + modifiedSentence); clientSocket.close();

FTP: the file transfer protocol



- transfer file to/from remote host
- client/server model
 - client: side that initiates transfer (either to/from remote)
 - server: remote host
- * ftp: RFC 959
- ftp server: port 21

Application Layer 2-3

FTP: separate control, data connections

- FTP client contacts FTP server at port 21, using TCP
- client authorized over control connection
- client browses remote directory, sends commands over control connection
- when server receives file transfer command, server opens 2nd TCP data connection (for file) to client
- after transferring one file, server closes data connection



- server opens another TCP data connection to transfer another file
- control connection: "out of band"
- FTP server maintains "state": current directory, earlier authentication

FTP commands, responses

sample commands:

- sent as ASCII text over control channel
- ✤ USER username
- * PASS password
- LIST return list of file in current directory
- RETR filename retrieves (gets) file
- STOR filename stores (puts) file onto remote host

sample return codes

- status code and phrase (as in HTTP)
- 331 Username OK, password required
- 125 data connection already open; transfer starting
- 425 Can't open data connection
- 452 Error writing
 file

Application Layer 2-5

Electronic mail

Three major components:

- user agents
- mail servers
- simple mail transfer protocol: SMTP

User Agent

- * a.k.a. "mail reader"
- composing, editing, reading mail messages
- e.g., Outlook, Thunderbird, iPhone mail client
- outgoing, incoming messages stored on server



Electronic mail: mail servers

mail servers:

- mailbox contains incoming messages for user
- message queue of outgoing (to be sent) mail messages
- SMTP protocol between mail servers to send email messages
 - client: sending mail server
 - "server": receiving mail server



Electronic Mail: SMTP [RFC 2821]

- uses TCP to reliably transfer email message from client to server, port 25
- direct transfer: sending server to receiving server
- three phases of transfer
 - handshaking (greeting)
 - transfer of messages
 - closure
- command/response interaction (like HTTP, FTP)
 - commands: ASCII text
 - response: status code and phrase
- messages must be in 7-bit ASCI

Scenario: Alice sends message to Bob

- I) Alice uses UA to compose message "to" bob@someschool.edu
- 2) Alice's UA sends message to her mail server; message placed in message queue
- client side of SMTP opens TCP connection with Bob's mail server
- SMTP client sends Alice's message over the TCP connection
- 5) Bob's mail server places the message in Bob's mailbox
- 6) Bob invokes his user agent to read message



Application Layer 2-9

Sample SMTP interaction

S: 220 hamburger.edu C: HELO crepes.fr S: 250 Hello crepes.fr, pleased to meet you C: MAIL FROM: <alice@crepes.fr> S: 250 alice@crepes.fr... Sender ok C: RCPT TO: <bob@hamburger.edu> S: 250 bob@hamburger.edu ... Recipient ok C: DATA S: 354 Enter mail, end with "." on a line by itself C: Do you like ketchup? C: How about pickles? C: . S: 250 Message accepted for delivery C: QUIT S: 221 hamburger.edu closing connection

Try SMTP interaction for yourself:

- * telnet servername 25
- see 220 reply from server
- enter HELO, MAIL FROM, RCPT TO, DATA, QUIT commands

above lets you send email without using email client (reader)

Application Layer 2-11

SMTP: final words

- SMTP uses persistent connections
- SMTP requires message (header & body) to be in 7-bit ASCII
- SMTP server uses CRLF.CRLF to determine end of message

comparison with HTTP:

- HTTP: pull
- SMTP: push
- both have ASCII command/response interaction, status codes
- HTTP: each object encapsulated in its own response msg
- SMTP: multiple objects sent in multipart msg

Mail message format



Application Layer 2-13

Mail access protocols



- SMTP: delivery/storage to receiver's server
- mail access protocol: retrieval from server
 - POP: Post Office Protocol [RFC 1939]: authorization, download
 - IMAP: Internet Mail Access Protocol [RFC 1730]: more features, including manipulation of stored msgs on server
 - HTTP: gmail, Hotmail, Yahoo! Mail, etc.



Chapter 2: outline

2.5 DNS

DNS: domain name system

people: many identifiers:

- SSN, name, passport # Internet hosts, routers:
 - IP address (32 bit) used for addressing datagrams
 - "name", e.g., www.yahoo.com used by humans
- <u>Q:</u> how to map between IP address and name, and vice versa ?

Domain Name System:

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

Application Layer 2-17

DNS: services, structure

DNS services

- hostname to IP address translation
- host aliasing
 - canonical, alias names
- mail server aliasing
- load distribution
 - replicated Web servers: many IP addresses correspond to one name

why not centralize DNS?

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

A: doesn't scale!

DNS: a distributed, hierarchical database



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

- client queries 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

Application Layer 2-19

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



TLD, authoritative servers

top-level domain (TLD) servers:

- responsible for com, org, net, edu, aero, jobs, museums, and all top-level country domains, e.g.: uk, fr, ca, jp
- Network Solutions maintains servers for .com TLD
- Educause for .edu TLD

authoritative DNS servers:

- organization's own DNS server(s), providing authoritative hostname to IP mappings for organization's named hosts
- can be maintained by organization or service provider

Application Layer 2-21

Local DNS name server

- does not strictly belong to hierarchy
- each ISP (residential ISP, company, university) has one
 - also called "default name server"
- when host makes DNS query, query is sent to its local DNS server
 - has local cache of recent name-to-address translation pairs (but may be out of date!)
 - acts as proxy, forwards query into hierarchy



 host at cis.poly.edu wants IP address for gaia.cs.umass.edu

iterated query:

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



Application Layer 2-23





DNS: caching, updating records

- once (any) name server learns mapping, it caches mapping
 - cache entries timeout (disappear) after some time (TTL)
 - TLD servers typically cached in local name servers
 thus root name servers not often visited
- cached entries may be out-of-date (best effort name-to-address translation!)
 - if name host changes IP address, may not be known Internet-wide until all TTLs expire
- update/notify mechanisms proposed IETF standard
 - RFC 2136

Application Layer 2-25

DNS records

DNS: distributed db storing resource records (RR)

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

type=A

- name is hostname
- value is IP address

<u>type=NS</u>

- name is domain (e.g., foo.com)
- value is hostname 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

<u>type=MX</u>

 value is name of mailserver associated with name

Explore DNS

- Use nslookup to find IP addresses
- Use whois to find domain registration

Application Layer 2-27

Attacking DNS

DDoS attacks

- Bombard root servers with traffic
 - Not successful to date
 - Traffic Filtering
 - Local DNS servers cache IPs of TLD servers, allowing root server bypass
- Bombard TLD servers
 - Potentially more dangerous

Redirect attacks

- Man-in-middle
 - Intercept queries
- DNS poisoning
 - Send bogus replies to DNS server, which caches

Exploit DNS for DDoS

- Send queries with spoofed source address: target IP
- Requires amplification

Chapter 2: outline

2.6 P2P applications

Application Layer 2-29

Pure P2P architecture

- ✤ no always-on server
- arbitrary end systems directly communicate
- peers are intermittently connected and change IP addresses

examples:

- file distribution (BitTorrent)
- Streaming (KanKan)
- VoIP (Skype)



File distribution: client-server vs P2P

<u>Question</u>: how much time to distribute file (size F) from one server to N peers?

peer upload/download capacity is limited resource



Application Layer 2-31

File distribution time: client-server



File distribution time: P2P

- server transmission: must upload at least one copy
 - time to send one copy: F/u_s
- client: each client must download file copy
 min client download time: F/d_{min}



- clients: as aggregate must download NF bits
 - max upload rate (limiting max download rate) is $u_s + \Sigma u_i$



increases linearly in \hat{N} ...

Application Layer 2-33

Client-server vs. P2P: example

client upload rate = u, F/u = 1 hour, $u_s = 10u$, $d_{min} \ge u_s$



P2P file distribution: BitTorrent

- file divided into 256Kb chunks
- peers in torrent send/receive file chunks



P2P file distribution: BitTorrent

- peer joining torrent:
 - has no chunks, but will accumulate them over time from other peers
 - registers with tracker to get list of peers, connects to subset of peers ("neighbors")



- while downloading, peer uploads chunks to other peers
- peer may change peers with whom it exchanges chunks
- churn: peers may come and go
- once peer has entire file, it may (selfishly) leave or (altruistically) remain in torrent

BitTorrent: requesting, sending file chunks

requesting chunks:

- at any given time, different peers have different subsets of file chunks
- periodically, Alice asks each peer for list of chunks that they have
- Alice requests missing chunks from peers, rarest first

sending chunks: tit-for-tat

- Alice sends chunks to those four peers currently sending her chunks at highest rate
 - other peers are choked by Alice (do not receive chunks from her)
 - re-evaluate top 4 every10 secs
- every 30 secs: randomly select another peer, starts sending chunks
 - "optimistically unchoke" this peer
 - newly chosen peer may join top 4

Application Layer 2-37

BitTorrent: tit-for-tat

- (1) Alice "optimistically unchokes" Bob
- (2) Alice becomes one of Bob's top-four providers; Bob reciprocates
- (3) Bob becomes one of Alice's top-four providers



Distributed Hash Table (DHT)

- DHT: a distributed P2P database
- * database has (key, value) pairs; examples:
 - key: ss number; value: human name
 - key: movie title; value: IP address
- Distribute the (key, value) pairs over the (millions of peers)
- * a peer queries DHT with key
 - DHT returns values that match the key
- peers can also insert (key, value) pairs

Application 2-39

Q: how to assign keys to peers?

- central issue:
 - assigning (key, value) pairs to peers.
- ✤ basic idea:
 - convert each key to an integer
 - Assign integer to each peer
 - put (key,value) pair in the peer that is closest to the key

DHT identifiers

- assign integer identifier to each peer in range
 [0,2ⁿ-1] for some n.
 - each identifier represented by *n* bits.
- require each key to be an integer in same range
- to get integer key, hash original key
 - e.g., key = hash("Led Zeppelin IV")
 - this is why its is referred to as a distributed "hash" table

Application 2-41

Assign keys to peers

- rule: assign key to the peer that has the closest ID.
- convention in lecture: closest is the immediate successor of the key.
- ✤ e.g., *n*=4; peers: 1,3,4,5,8,10,12,14;
 - key = 13, then successor peer = 14
 - key = 15, then successor peer = 1

Application 2-42



- each peer *only* aware of immediate successor and predecessor.
- overlay network"

Application 2-43

Circular DHT (I)



Application 2-44

Circular DHT with shortcuts



- each peer keeps track of IP addresses of predecessor, successor, short cuts.
- reduced from 6 to 2 messages.
- possible to design shortcuts so O(log N) neighbors, O(log N) messages in query

Application 2-45



handling peer churn:

peers may come and go (churn)
 each peer knows address of its
 two successors

*each peer periodically pings its two successors to check aliveness

 if immediate successor leaves, choose next successor as new immediate successor

example: peer 5 abruptly leaves

peer 4 detects peer 5 departure; makes 8 its immediate successor; asks 8 who its immediate successor is; makes 8' s immediate successor its second successor.
what if peer 13 wants to join?

Application 2-46

Chapter 2: summary

our study of network apps now complete!

- application architectures
 - client-server
 - P2P
- application service requirements:
 - reliability, bandwidth, delay
- Internet transport service model
 - connection-oriented, reliable: TCP
 - unreliable, datagrams: UDP

- specific protocols:
 - HTTP
 - FTP
 - SMTP
 - DNS
 - P2P: BitTorrent, DHT
- socket programming: TCP, UDP sockets

Application Layer 2-47

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

important themes:

- control vs. data msgs
 - in-band, out-of-band
- centralized vs. decentralized
- stateless vs. stateful
- reliable vs. unreliable msg transfer
- "complexity at network edge"