

Section 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 applications
- 2.7 Socket programming with UDP
- 2.8 Socket programming with TCP

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"

DNS

DNS services


- ❑ hostname to IP address translation
- ❑ host aliasing
 - ❖ Canonical, 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!*

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

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 and Authoritative Servers

- ❑ **Top-level domain (TLD) servers:**
 - ❖ responsible for com, org, net, edu, etc, and all top-level country domains uk, fr, ca, jp.
 - ❖ 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, mail).
 - ❖ can be maintained by organization or service provider

Local 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
 - ❖ acts as proxy, forwards query into hierarchy


2: Application Layer 68

DNS name resolution example

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

iterated query:

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




2: Application Layer 69

DNS name resolution example

recursive query:

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



2: Application Layer 70

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>

2: Application Layer 71

DNS records

DNS: distributed db storing resource records (RR)

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

- r Type=A
 - ❖ name is hostname
 - ❖ value is IP address
- ❑ Type=NS
 - ❖ name is domain (e.g. foo.com)
 - ❖ value is hostname of authoritative name server for this domain
- r 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
- r Type=MX
 - ❖ value is name of mailserver associated with name

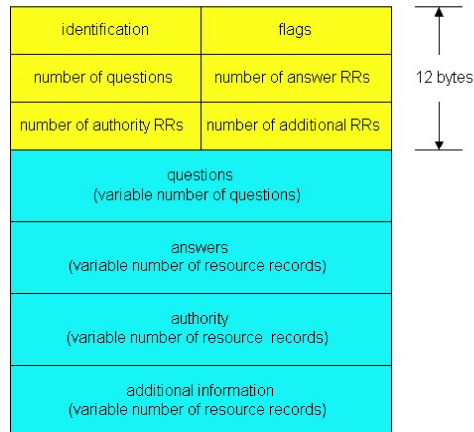
2: Application Layer 72

DNS protocol, messages

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

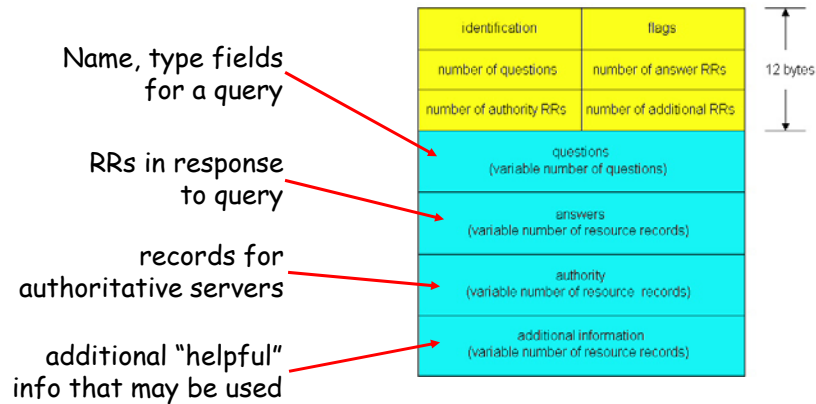
msg header

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



2: Application Layer 73

DNS protocol, messages



2: Application Layer 74

Inserting records into DNS

- ❑ example: new startup "Network Utopia"
- ❑ register name networkutopia.com at *DNS registrar* (e.g., Network Solutions)
 - ❖ provide names, IP addresses of authoritative name server (primary and secondary)
 - ❖ registrar inserts two RRs into com TLD server:

(networkutopia.com, dns1.networkutopia.com, NS)
(dns1.networkutopia.com, 212.212.212.1, A)
- ❑ create authoritative server Type A record for www.networkutopia.com; Type MX record for networkutopia.com
- ❑ *How do people get IP address of your Web site?*

2: Application Layer 75


Section 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 applications**
- ❑ 2.7 Socket programming with UDP
- ❑ 2.8 Socket programming with TCP

2: Application Layer 76

Pure P2P architecture


- ❑ *no* always-on server
- ❑ arbitrary end systems directly communicate
- ❑ peers are intermittently connected and change IP addresses
- ❑ **Three topics:**
 - ❖ File distribution
 - ❖ Searching for information
 - ❖ Case Study: Skype



2: Application Layer 77


File Distribution: Server-Client vs P2P

Question: How much time to distribute file from one server to N peers?



File distribution time: server-client

- server sequentially sends N copies:
 - ❖ NF/u_s time
- client i takes F/d_i time to download




Time to distribute F to N clients using client/server approach = $d_{cs} = \max \{ NF/u_s, F/\min_i(d_i) \}$

increases linearly in N (for large N)

File distribution time: P2P

- server must send one copy: F/u_s time
- client i takes F/d_i time to download
- NF bits must be downloaded (aggregate)
 - fastest possible upload rate: $u_s + \sum u_i$




$$d_{P2P} = \max \left\{ F/u_s, F/\min(d_i), NF/(u_s + \sum u_i) \right\}$$

2: Application Layer 80

Server-client vs. P2P: example

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




2: Application Layer 81

File distribution: BitTorrent

r P2P file distribution

tracker: tracks peers participating in torrent


torrent: group of peers exchanging chunks of a file



2: Application Layer 82

BitTorrent (1)

- ❑ file divided into 256KB *chunks*.
- ❑ peer joining torrent:
 - ❖ has no chunks, but will accumulate them over time
 - ❖ registers with tracker to get list of peers, connects to subset of peers ("neighbors")
- ❑ while downloading, peer uploads chunks to other peers.
- ❑ peers may come and go
- ❑ once peer has entire file, it may (selfishly) leave or (altruistically) remain



2: Application Layer 83

BitTorrent (2)

Pulling Chunks

- ❑ at any given time, different peers have different subsets of file chunks
- ❑ periodically, a peer (Alice) asks each neighbor for list of chunks that they have.
- ❑ Alice sends requests for her missing chunks
 - ❖ rarest first


Sending Chunks: tit-for-tat

- r Alice sends chunks to four neighbors currently sending her chunks *at the highest rate*
 - ❖ re-evaluate top 4 every 10 secs
- r every 30 secs: randomly select another peer, starts sending chunks
 - ❖ newly chosen peer may join top 4
 - ❖ "optimistically unchoke"

2: Application Layer 84

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



2: Application Layer 85

Distributed Hash Table (DHT)

- DHT = distributed P2P database
- Database has (key, value) pairs;
 - ❖ key: ss number; value: human name
 - ❖ key: content type; value: IP address
- Peers query DB with key
 - ❖ DB returns values that match the key
- Peers can also insert (key, value) peers


DHT Identifiers

- Assign integer identifier to each peer in range $[0, 2^n - 1]$.
 - ❖ Each identifier can be represented by n bits.
- Require each key to be an integer in same range.
- To get integer keys, hash original key.
 - ❖ eg, key = h("Led Zeppelin IV")
 - ❖ This is why they call it a distributed "hash" table

How to assign keys to peers?

- Central issue:
 - ❖ Assigning (key, value) pairs to peers.
- Rule: assign key to the peer that has the **closest** ID.
- Convention in lecture: closest is the **immediate successor** of the key.
- Ex: $n=4$; peers: 1,3,4,5,8,10,12,14;
 - ❖ key = 13, then successor peer = 14
 - ❖ key = 15, then successor peer = 1


Circular DHT (1)



- Each peer *only* aware of immediate successor and predecessor.
- "Overlay network"


Circle DHT (2)

$O(N)$ messages
on avg to resolve
query, when there
are N peers




Define closest
as closest
successor

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

Peer Churn




- To handle peer churn, require each peer to know the IP address of its two successors.
- Each peer periodically pings its two successors to see if they are still alive.

- ❑ Peer 5 abruptly leaves
- ❑ Peer 4 detects; 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?


P2P Case study: Skype

- ❑ inherently P2P: pairs of users communicate.
- ❑ proprietary application-layer protocol (inferred via reverse engineering)
- ❑ hierarchical overlay with SNs
- ❑ Index maps usernames to IP addresses; distributed over SNs



Peers as relays

- Problem when both Alice and Bob are behind "NATs".
 - ❖ NAT prevents an outside peer from initiating a call to insider peer
- Solution:
 - ❖ Using Alice's and Bob's SNs, Relay is chosen
 - ❖ Each peer initiates session with relay.
 - ❖ Peers can now communicate through NATs via relay



2: Application Layer 94

Section 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 applications
- 2.7 Socket programming with UDP
- 2.8 Socket programming with TCP

2: Application Layer 95

Socket programming

Goal: learn how to build client/server application that communicate using sockets

Socket API

- ❑ introduced in BSD4.1 UNIX, 1981
- ❑ explicitly created, used, released by apps
- ❑ client/server paradigm
- ❑ two types of transport service via socket API:
 - ❖ UDP
 - ❖ TCP

socket

A *application-created, OS-controlled* interface (a "door") into which application process can **both send and receive** messages to/from another application process

2: Application Layer 96

Socket programming basics

- ❑ Server must be **running** before client can send anything to it.
- ❑ Server must have a **socket** (door) through which it receives and sends segments
- ❑ Similarly client needs a socket
- ❑ Socket is locally identified with a **port number**
 - ❖ Analogous to the apt # in a building
- ❑ Client **needs to know** server IP address and socket port number.

2: Application Layer 97

Socket programming *with UDP*

UDP: no "connection" between client and server

- ❑ no handshaking
- ❑ sender explicitly attaches IP address and port of destination to each segment
- ❑ OS attaches IP address and port of sending socket to each segment
- ❑ Server can extract IP address, port of sender from received segment

application viewpoint

UDP provides unreliable transfer of groups of bytes ("datagrams") between client and server

Note: the official terminology for a UDP packet is "datagram". In this class, we instead use "UDP segment".

2: Application Layer 98

Client/server socket interaction: UDP

Server

create socket,
port= x.
serverSocket =
DatagramSocket()

read datagram from
serverSocket

write reply to
serverSocket
specifying
client address,
port number

Client

create socket,
clientSocket =
DatagramSocket()

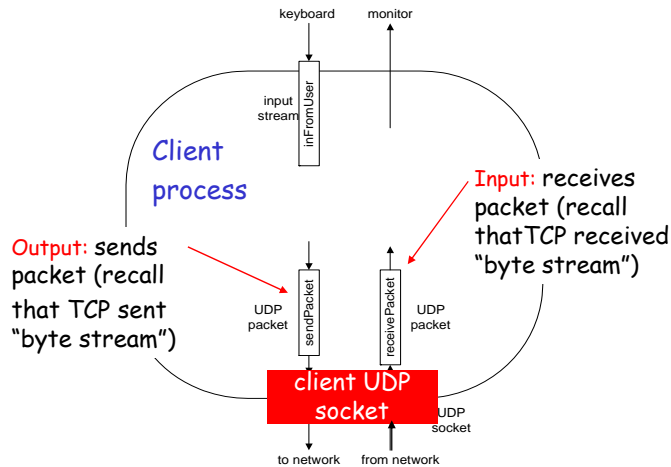
Create datagram with server IP and
port=x; send datagram via
clientSocket

read datagram from
clientSocket

close
clientSocket

2: Application Layer 99

Example: Java client (UDP)



2: Application Layer 100

UDP observations & questions

- ❑ Both client server use DatagramSocket
- ❑ Dest IP and port are explicitly attached to segment.
- ❑ What would happen if change both clientSocket and serverSocket to "mySocket"?
- ❑ Can the client send a segment to server without knowing the server's IP address and/or port number?
- ❑ Can multiple clients use the server?

2: Application Layer 101


Section 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 applications
- ❑ 2.7 Socket programming with UDP
- ❑ 2.8 Socket programming with TCP

2: Application Layer 102

Socket-programming using TCP

TCP service: reliable transfer of **bytes** from one process to another



2: Application Layer 103

Socket programming *with TCP*

Client must contact server

- server process must first be running
- server must have created socket (door) that welcomes client's contact

Client contacts server by:

- creating client-local TCP socket
- specifying IP address, port number of server process
- When **client creates socket**: client TCP establishes connection to server TCP

- When contacted by client, **server TCP creates new socket** for server process to communicate with client
 - allows server to talk with multiple clients
 - source port numbers used to distinguish clients (more in Chap 3)

application viewpoint

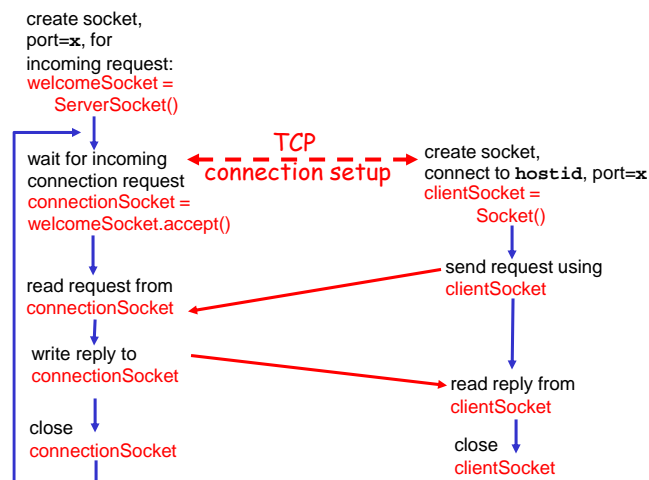
TCP provides reliable, in-order transfer of bytes ("pipe") between client and server

2: Application Layer 104

Client/server socket interaction: TCP

Server (running on `hostid`)


Client



2: Application Layer 105

Stream jargon

- ❑ A **stream** is a sequence of characters that flow into or out of a process.
- ❑ An **input stream** is attached to some input source for the process, e.g., keyboard or socket.
- ❑ An **output stream** is attached to an output source, e.g., monitor or socket.



2: Application Layer 106

TCP observations & questions

- ❑ Server has two types of sockets:
 - ❖ ServerSocket and Socket
- ❑ When client knocks on serverSocket's "door," server creates connectionSocket and completes TCP conx.
- ❑ Dest IP and port are **not** explicitly attached to segment.
- ❑ Can **multiple clients** use the server?

2: Application Layer 107

CS2505 labs

- The labs complement this section of the course by allowing you to write your own simple client/sever programs using
 - ❖ Java
 - ❖ UDP
 - ❖ TCP

Section 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
- r specific protocols:
 - ❖ HTTP
 - ❖ FTP
 - ❖ SMTP, POP, IMAP
 - ❖ DNS
 - ❖ P2P: BitTorrent, Skype
- r socket programming

Section 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:

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

2: Application Layer 110