

# Last time

- TCP
  - ◆ Throughput
  - ◆ Fairness
  - ◆ Delay modeling
- TCP socket programming


# This time

- NAT
  
- Application layer
  - ◆ Intro
  
  - ◆ Web / HTTP

# Chapter 4: Network Layer

- 4. 1 Introduction
- 4.2 Virtual circuit and datagram networks
- 4.3 What's inside a router
- **4.4 IP: Internet Protocol**
  - ◆ Datagram format
  - ◆ **IPv4 addressing**
  - ◆ ICMP
  - ◆ IPv6
- 4.5 Routing algorithms
  - ◆ Link state
  - ◆ Distance Vector
  - ◆ Hierarchical routing
- 4.6 Routing in the Internet
  - ◆ RIP
  - ◆ OSPF
  - ◆ BGP
- 4.7 Broadcast and multicast routing

# NAT: Network Address Translation



*All* datagrams *leaving* local network have **same** single source NAT IP address: 138.76.29.7, different source port numbers

Datagrams with source or destination in this network have 10.0.0/24 address for source, destination (as usual)

# NAT: Network Address Translation

- **Motivation:** local network uses just one IP address as far as outside world is concerned:
  - ◆ range of addresses not needed from ISP: just one IP address for all devices
  - ◆ can change addresses of devices in local network without notifying outside world
  - ◆ can change ISP without changing addresses of devices in local network
  - ◆ devices inside local net not explicitly addressable, visible by outside world (a security plus).

# NAT: Network Address Translation

**Implementation:** NAT router must:

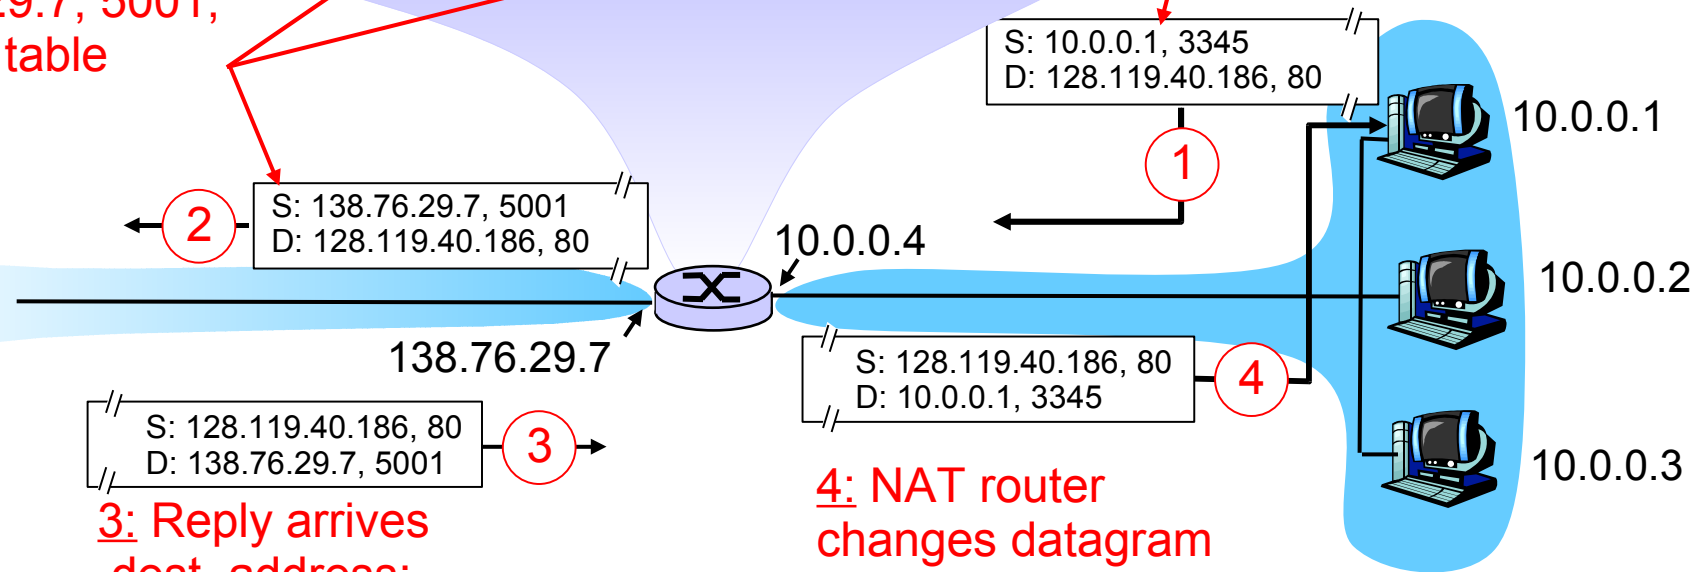
- ◆ *outgoing datagrams: replace* (source IP address, port #) of every outgoing datagram to (NAT IP address, new port #)
  - . . . remote clients/servers will respond using (NAT IP address, new port #) as destination addr.
- ◆ *remember (in NAT translation table)* every (source IP address, port #) to (NAT IP address, new port #) translation pair
- ◆ *incoming datagrams: replace* (NAT IP address, new port #) in dest fields of every incoming datagram with corresponding (source IP address, port #) stored in NAT table

# NAT: Network Address Translation

NAT translation table	
WAN side addr	LAN side addr
138.76.29.7, 5001	10.0.0.1, 3345
.....	.....

1: host 10.0.0.1 sends datagram to 128.119.40.186, 80

2: NAT router changes datagram source addr from 10.0.0.1, 3345 to 138.76.29.7, 5001, updates table



3: Reply arrives  
dest. address:  
138.76.29.7, 5001

4: NAT router changes datagram dest addr from 138.76.29.7, 5001 to 10.0.0.1, 3345

# NAT: Network Address Translation

- 16-bit port-number field:
  - ◆ 60,000 simultaneous connections with a single LAN-side address!
- NAT is controversial:
  - ◆ routers should only process up to layer 3
  - ◆ violates end-to-end argument
    - NAT possibility must be taken into account by app designers, eg, P2P applications
  - ◆ address shortage should instead be solved by IPv6



# 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

# Some network apps

- E-mail
- Web
- Instant messaging
- Remote login
- P2P file sharing
- Multi-user network games
- Streaming stored video clips
- Internet telephone
- Real-time video conference
- Massive parallel computing


# Creating a network app

## Write programs that

- ◆ run on different end systems and
- ◆ communicate over a network.
- ◆ e.g., Web: Web server software communicates with browser software

## Little software written for devices in network core

- ◆ network core devices do not run user application code
- ◆ application on end systems allows for rapid app development, propagation




# 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

# Application architectures

- Client-server
- Peer-to-peer (P2P)
- Hybrid of client-server and P2P

# Client-server architecture



## server:

- ◆ always-on host
- ◆ permanent IP address
- ◆ server farms for scaling


## clients:

- ◆ communicate with server
- ◆ may be intermittently connected
- ◆ may have dynamic IP addresses
- ◆ do not communicate directly with each other

# Pure P2P architecture

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

Highly scalable but difficult to manage



# Hybrid of client-server and P2P

## Skype

- ◆ Internet telephony app
- ◆ Finding address of remote party: centralized server(s)
- ◆ Client-client connection is direct (not through server)

## Instant messaging

- ◆ Chatting between two users can be P2P
- ◆ Presence detection/location centralized:
  - User registers its IP address with central server when it comes online
  - User contacts central server to find IP addresses of buddies



# Processes communicating

**Process:** program running within a host.

- within same host, two processes communicate using **inter-process communication** (defined by OS).
- processes in different hosts communicate by exchanging **messages**

**Client process:** process that initiates communication

**Server process:** process that waits to be contacted

- Note: applications with P2P architectures have client processes & server processes

# App-layer protocol defines

- Types of messages exchanged,
  - ◆ e.g., request, response
- Message syntax:
  - ◆ what fields in messages & how fields are delineated
- Message semantics
  - ◆ meaning of information in fields
- Rules for when and how processes send & respond to messages

## Public-domain protocols:

- defined in RFCs
- allows for interoperability
- e.g., HTTP, SMTP

## Proprietary protocols:

- e.g., KaZaA

# What transport service does an app need?

## Data loss

- some apps (e.g., audio) can tolerate some loss
- other apps (e.g., file transfer, telnet) require 100% reliable data transfer

## Timing

- some apps (e.g., Internet telephony, interactive games) require low delay to be “effective”

## Bandwidth

- some apps (e.g., multimedia) require minimum amount of bandwidth to be “effective”
- other apps (“elastic apps”) make use of whatever bandwidth they get

## Transport service requirements of common apps

<b>Application</b>	<b>Data loss</b>	<b>Bandwidth</b>	<b>Time Sensitive</b>
file transfer	no loss	elastic	no
e-mail	no loss	elastic	no
Web documents	no loss	elastic	no
real-time audio/video	loss-tolerant	audio: 5kbps-1Mbps video: 10kbps-5Mbps	yes, 100's msec
stored audio/video	loss-tolerant	same as above	yes, few secs
interactive games	loss-tolerant	few kbps up	yes, 100's msec
instant messaging	no loss	elastic	yes and no

# Internet transport protocols services

## TCP service:

- *connection-oriented*: setup required between client and server processes
- *reliable transport* between sending and receiving process
- *flow control*: sender won't overwhelm receiver
- *congestion control*: throttle sender when network overloaded
- *does not provide*: timing, minimum bandwidth guarantees

## UDP service:

- unreliable data transfer between sending and receiving process
- does not provide: connection setup, reliability, flow control, congestion control, timing, or bandwidth guarantee

Q: why bother? Why is there a UDP?

## Internet apps: application, transport protocols

<b>Application</b>	<b>Application layer protocol</b>	<b>Underlying transport protocol</b>
e-mail	SMTP [RFC 2821]	TCP
remote terminal access	Telnet [RFC 854]	TCP
Web	HTTP [RFC 2616]	TCP
file transfer	FTP [RFC 959]	TCP
streaming multimedia	proprietary (e.g. RealNetworks)	TCP or UDP
Internet telephony	proprietary (e.g., Vonage, Dialpad)	typically UDP

# 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
- 2.7 Socket programming with TCP
- 2.8 Socket programming with UDP
- 2.9 Building a Web server

# Web and HTTP

## First some jargon

- **Web page** consists of **objects**
- Object can be HTML file, JPEG image, Java applet, audio file,...
- Web page consists of **base HTML-file** which includes several referenced objects
- Each object is addressable by a **URL**
- Example URL:

`http://www.someschool.edu/someDept/pic.gif`

host name


path name



# HTTP overview

## HTTP: hypertext transfer protocol

- Web's application layer protocol
- client/server model
  - ◆ *client*: browser that requests, receives, "displays" Web objects
  - ◆ *server*: Web server sends objects in response to requests
- HTTP 1.0: RFC 1945
- HTTP 1.1: RFC 2068



# HTTP overview (continued)

## Uses TCP:

- client initiates TCP connection (creates socket) to server, port 80
- server accepts TCP connection from client
- HTTP messages (application-layer protocol messages) exchanged between browser (HTTP client) and Web server (HTTP server)
- TCP connection closed

## HTTP is “stateless”

- server maintains no information about past client requests

### **Protocols that maintain “state” are complex!** aside

- past history (state) must be maintained
- if server/client crashes, their views of “state” may be inconsistent, must be reconciled

# HTTP connections

## Nonpersistent HTTP

- At most one object is sent over a TCP connection.
- HTTP/1.0 uses nonpersistent HTTP

## Persistent HTTP

- Multiple objects can be sent over single TCP connection between client and server.
- HTTP/1.1 uses persistent connections in default mode

# Nonpersistent HTTP

Suppose user enters URL

`http://www.someSchool.edu/someDept/home.index` (contains text, references to 10 jpeg images)

1a. HTTP client initiates TCP connection to HTTP server (process) at `www.someSchool.edu` on port 80

1b. HTTP server at host `www.someSchool.edu` waiting for TCP connection at port 80. "accepts" connection, notifying client

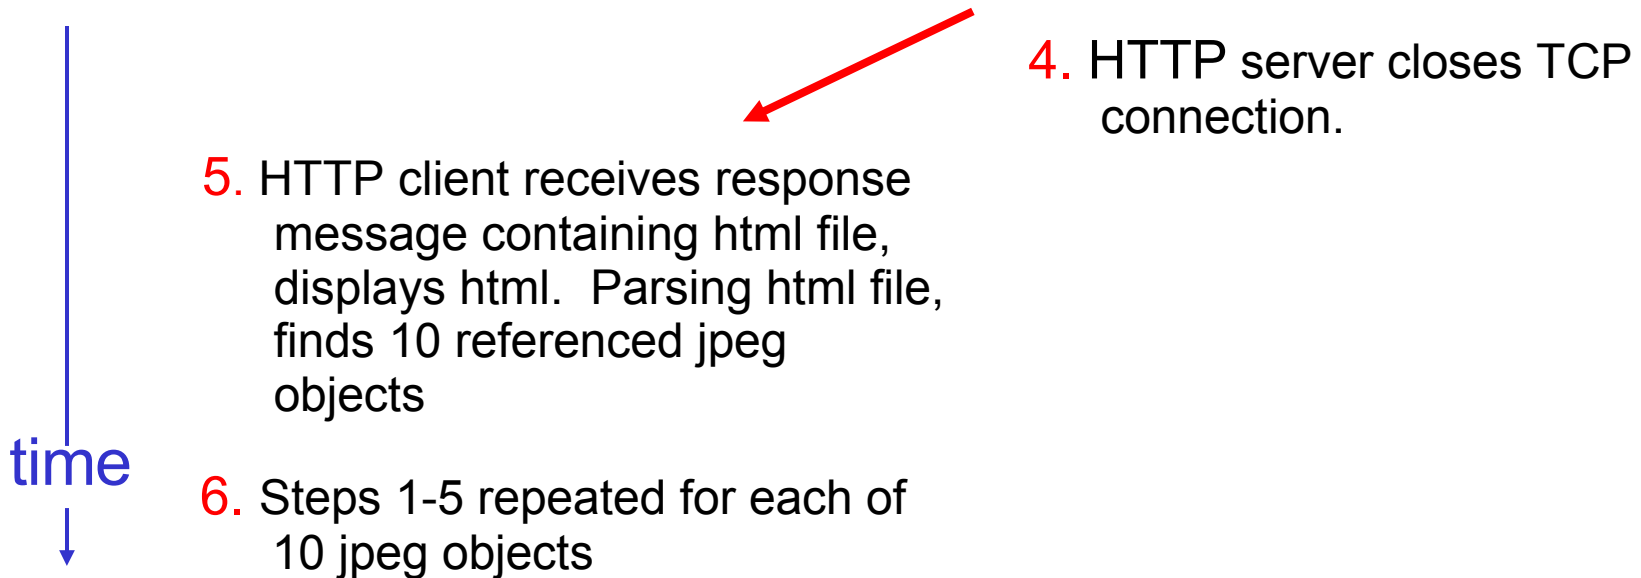
2. HTTP client sends HTTP *request message* (containing URL) into TCP connection socket. Message indicates that client wants object `/someDept/home.index`

3. HTTP server receives request message, forms *response message* containing requested object, and sends message into its socket

time



# Nonpersistent HTTP (cont.)




# Non-Persistent HTTP: Response time

**Definition of RTT:** time to send a small packet to travel from client to server and back.

## Response time:

- one RTT to initiate TCP connection
- one RTT for HTTP request and first few bytes of HTTP response to return
- file transmission time

**total =  $2RTT + \text{transmit time}$**



# Persistent HTTP

## Nonpersistent HTTP issues:

- requires 2 RTTs per object
- OS overhead for *each* TCP connection
- 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 sent over open 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

# Recap

- NAT
  
- Application layer
  - ◆ Intro
  
  - ◆ Web / HTTP



# Next time

- Finish HTTP
- FTP
- SMTP (email)