


Chapter 2 Application Layer



Computer Networking: A Top-Down Approach
6th edition
Jim Kurose, Keith Ross
Addison-Wesley
March 2012

A note on the use of these ppt slides:
We're making these slides freely available to all (faculty, students, readers). They're in PowerPoint form so you can see the animations, and can add, modify, and delete slides (including this one) and slide content to suit your needs. They obviously represent a lot of work on our part. In return for use, we only ask the following:

- If you use these slides (e.g., in a class) that you mention their source (after all, we'd like people to use our book)
- If you post any slides on a www site, that you note that they are adapted from (or perhaps identical to) our slides, and note our copyright of this material.

Thanks and enjoy! JFK/KWR

© All material copyright 1996-2012
J.F. Kurose and K.W. Ross. All Rights Reserved

The course notes are adapted for Bucknell's CSCI 363
Xiannong Meng
Spring 2016

Application Layer 2-1

Application Layer 2-2

Chapter 2: outline

2.1 principles of network applications

- app architectures
- app requirement

2.6 P2P applications

2.2 Web and HTTP

2.3 FTP

2.4 electronic mail

- SMTP, POP3, IMAP

2.5 DNS

2.7 socket programming with UDP and TCP

Web and HTTP

First, a review...

- ❖ **web page** consists of **objects**
- ❖ object can be HTML file, JPEG image, Java script, audio file,...
- ❖ web page consists of **base HTML-file** which includes **several referenced objects**
- ❖ each object is addressable by a **URL**, e.g.,

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

protocol host name path name

Application Layer 2-3

HTTP overview

HTTP: hypertext transfer protocol

- ❖ Web's application layer protocol
- ❖ client/server model
 - **client:** browser that requests, receives, (using HTTP protocol) and "displays" Web objects
 - **server:** Web server sends (using HTTP protocol) objects in response to requests



Application Layer 2-4

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

aside
protocols that maintain "state" are complex!
❖ past history (state) must be maintained
❖ if server/client crashes, their views of "state" may be inconsistent, must be reconciled

Application Layer 2-5

HTTP connections

non-persistent HTTP

- ❖ at most one object sent over TCP connection
 - connection then closed
- ❖ downloading multiple objects required multiple connections


persistent HTTP

- ❖ multiple objects can be sent over single TCP connection between client, server


Application Layer 2-6

Non-persistent HTTP

suppose user enters URL: `www.someSchool.edu/someDepartment/home.index` (contains text, references to 10 jpeg images)



Non-persistent HTTP (cont.)




Non-persistent HTTP: response time

RTT (round-trip time): time for a small packet to travel from client to server and back

HTTP 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
- ❖ non-persistent HTTP response time = $2RTT + \text{file transmission time}$



Persistent HTTP

non-persistent HTTP issues:

- ❖ requires 2 RTTs per object
- ❖ OS overhead for each TCP connection
- ❖ browsers often open parallel TCP connections to fetch referenced objects
- ❖ HTTP 1.0

persistent HTTP:

- ❖ server leaves connection open after sending response
 - ❖ subsequent HTTP messages between same client/server sent over open connection
 - ❖ client sends requests as soon as it encounters a referenced object
 - ❖ as little as one RTT for all the referenced objects
 - ❖ HTTP 1.1
- Application Layer 2-10


HTTP request message

❖ two types of HTTP messages: *request*, *response*

- ❖ **HTTP request message:**
 - ASCII (human-readable format)



HTTP request message: general format



Uploading form input

POST method:

- ❖ web page often includes form input
- ❖ input is uploaded to server in entity body

URL method:

- ❖ uses GET method
- ❖ input is uploaded in URL field of request line:

```
www.somesite.com/animalsearch?monkeys&banana
```

Try it out with our own web server!

Application Layer 2-13

Method types

HTTP/1.0:


- ❖ GET
- ❖ POST
- ❖ HEAD
 - asks server to leave requested object out of response

HTTP/1.1:

- ❖ GET, POST, HEAD
- ❖ PUT
 - uploads file in entity body to path specified in URL field
- ❖ DELETE
 - deletes file specified in the URL field

Application Layer 2-14

HTTP response message



Application Layer 2-15

HTTP response status codes

- ❖ status code appears in 1st line in server-to-client response message.

- ❖ some sample codes:

- 200 OK**
 - request succeeded, requested object later in this msg
- 301 Moved Permanently**
 - requested object moved, new location specified later in this msg (Location:)
- 400 Bad Request**
 - request msg not understood by server
- 404 Not Found**
 - requested document not found on this server
- 505 HTTP Version Not Supported**

Application Layer 2-16

Trying out HTTP (client side) for yourself

1. Telnet to your favorite Web server:

```
telnet cis.poly.edu 80
```

opens TCP connection to port 80 (default HTTP server port) at cis.poly.edu. anything typed in sent to port 80 at cis.poly.edu

2. type in a GET HTTP request:

```
GET /~ross/ HTTP/1.1
Host: cis.poly.edu
```

by typing this in (hit carriage return twice), you send this minimal (but complete) GET request to HTTP server

3. look at response message sent by HTTP server!

(or use Wireshark to look at captured HTTP request/response)

Application Layer 2-17

User-server state: cookies

many Web sites use cookies
four components:


- 1) cookie header line of HTTP response message
- 2) cookie header line in next HTTP request message
- 3) cookie file kept on user's host, managed by user's browser
- 4) back-end database at Web site

example:

- ❖ Susan always access Internet from PC
- ❖ visits specific e-commerce site for first time
- ❖ when initial HTTP requests arrives at site, site creates:
 - unique ID
 - entry in backend database for ID

Application Layer 2-18

Cookies: keeping "state" (cont.)



Cookies (continued)

what cookies can be used for:

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

cookies and privacy: aside

- ❖ cookies permit sites to learn a lot about you
- ❖ you may supply name and e-mail to sites

how to keep "state":


- ❖ protocol endpoints: maintain state at sender/receiver over multiple transactions
- ❖ cookies: http messages carry state

Application Layer 2-20

Web caches (proxy server)

goal: satisfy client request without involving origin server
(proxy server can also by-pass censorship or any content control).

- ❖ 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



Application Layer 2-21

More about Web caching

- ❖ cache acts as both client and server
 - server for original requesting client
 - client to origin 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 (so too does P2P file sharing)

Application Layer 2-22


Caching example:

assumptions:

- ❖ avg object size: 100K bits
- ❖ avg request rate from browsers to origin servers: 15/sec
- ❖ avg data rate to browsers: 1.50 Mbps
- ❖ RTT from institutional router to any origin server: 2 sec
- ❖ access link rate: 1.54 Mbps

consequences:

- ❖ LAN utilization: 15%
- ❖ access link utilization = 97% **problem!**
- ❖ total delay = Internet delay + access delay + LAN delay = 2 sec + minutes + usecs



Application Layer 2-23


Caching example: fatter access link

assumptions:

- ❖ avg object size: 100K bits
- ❖ avg request rate from browsers to origin servers: 15/sec
- ❖ avg data rate to browsers: 1.50 Mbps
- ❖ RTT from institutional router to any origin server: 2 sec
- ❖ access link rate: 1.54 Mbps → 154 Mbps

consequences:

- ❖ LAN utilization: 15%
- ❖ access link utilization = 97% → 9.7%
- ❖ total delay = Internet delay + access delay + LAN delay = 2 sec + minutes + usecs



Application Layer 2-24

Cost: increased access link speed (not cheap!)

Caching example: install local cache

assumptions:


- ❖ avg object size: 100K bits
- ❖ avg request rate from browsers to origin servers: 15/sec
- ❖ avg data rate to browsers: 1.50 Mbps
- ❖ RTT from institutional router to any origin server: 2 sec
- ❖ access link rate: 1.54 Mbps

consequences:

- ❖ LAN utilization: 15%
- ❖ access link utilization = ?
- ❖ total delay = ?

How to compute link utilization, delay?

Cost: web cache (cheap!)




Application Layer 2-25

Caching example: install local cache

Calculating access link utilization, delay with cache:


- ❖ suppose cache hit rate is 0.4
 - 40% requests satisfied at cache, 60% requests satisfied at origin
- ❖ access link utilization:
 - 60% of requests use access link
- ❖ data rate to browsers over access link = $0.6 * 1.50 \text{ Mbps} = .9 \text{ Mbps}$
 - utilization = $0.9 / 1.54 = .58$
- ❖ total delay
 - = $0.6 * (\text{delay from origin servers}) + 0.4 * (\text{delay when satisfied at cache})$
 - = $0.6 (2.01) + 0.4 (0.01)$
 - = $\sim 1.2 \text{ secs}$
 - less than with 154 Mbps link (and cheaper too!)



Application Layer 2-26

Conditional GET

- ❖ **Goal:** don't send object if cache has up-to-date cached version
 - no object transmission delay
 - lower link utilization
- ❖ **cache:** specify date of cached copy in HTTP request
 - IF-modified-since: <date>
- ❖ **server:** response contains no object if cached copy is up-to-date:
 - HTTP/1.0 304 Not Modified



Application Layer 2-27