

# Chapter 2

## Application Layer


### 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 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



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

# 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

## **open protocols:**

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

## **proprietary protocols:**

- ❖ e.g., Skype

# What transport service does an app need?

## data integrity

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

## timing

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

## throughput

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

## security

- ❖ encryption, data integrity,  
...

# Internet transport protocols services

---

## TCP service:

- ❖ *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 throughput guarantee, security
- ❖ *connection-oriented*: setup required between client and server processes

## UDP service:

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

Q: why bother? Why is there a UDP?

# Chapter 2: outline

## 2.1 principles of network applications

- app architectures
- app requirements

## 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 and TCP

# Web and HTTP

*First, a review...*

- ❖ *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*, e.g.,

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

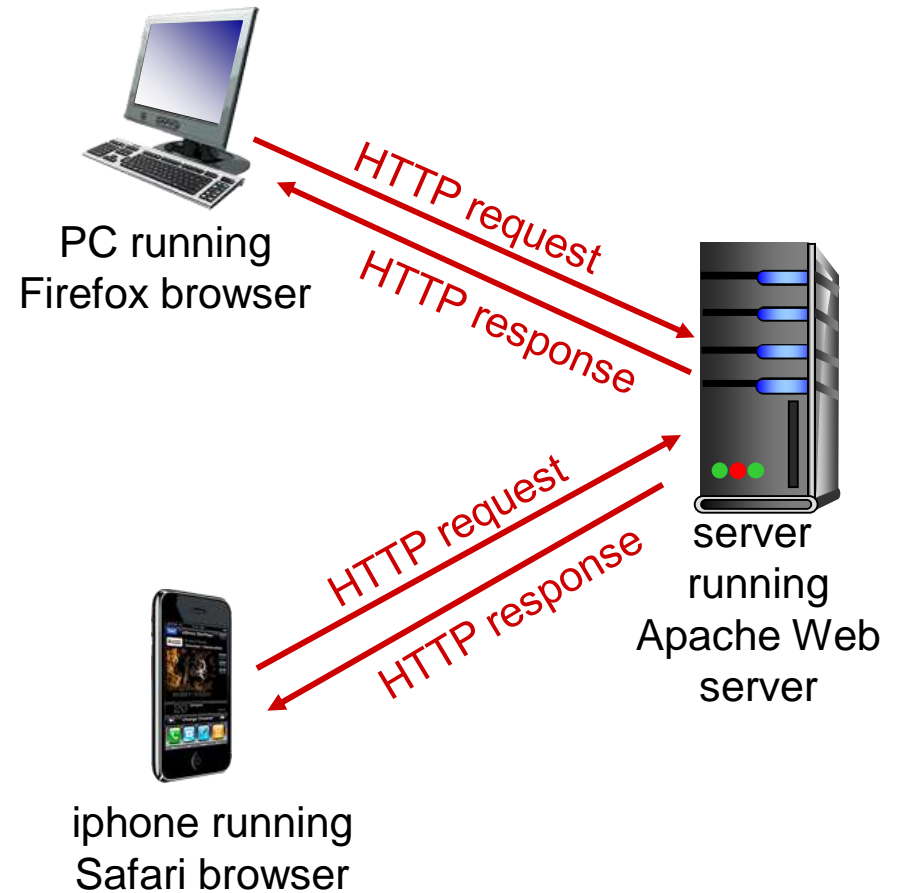
host name

path name

# HTTP overview

## HTTP: hypertext transfer protocol

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



# 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



# HTTP connections

## *non-persistent HTTP*

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

## *persistent HTTP*

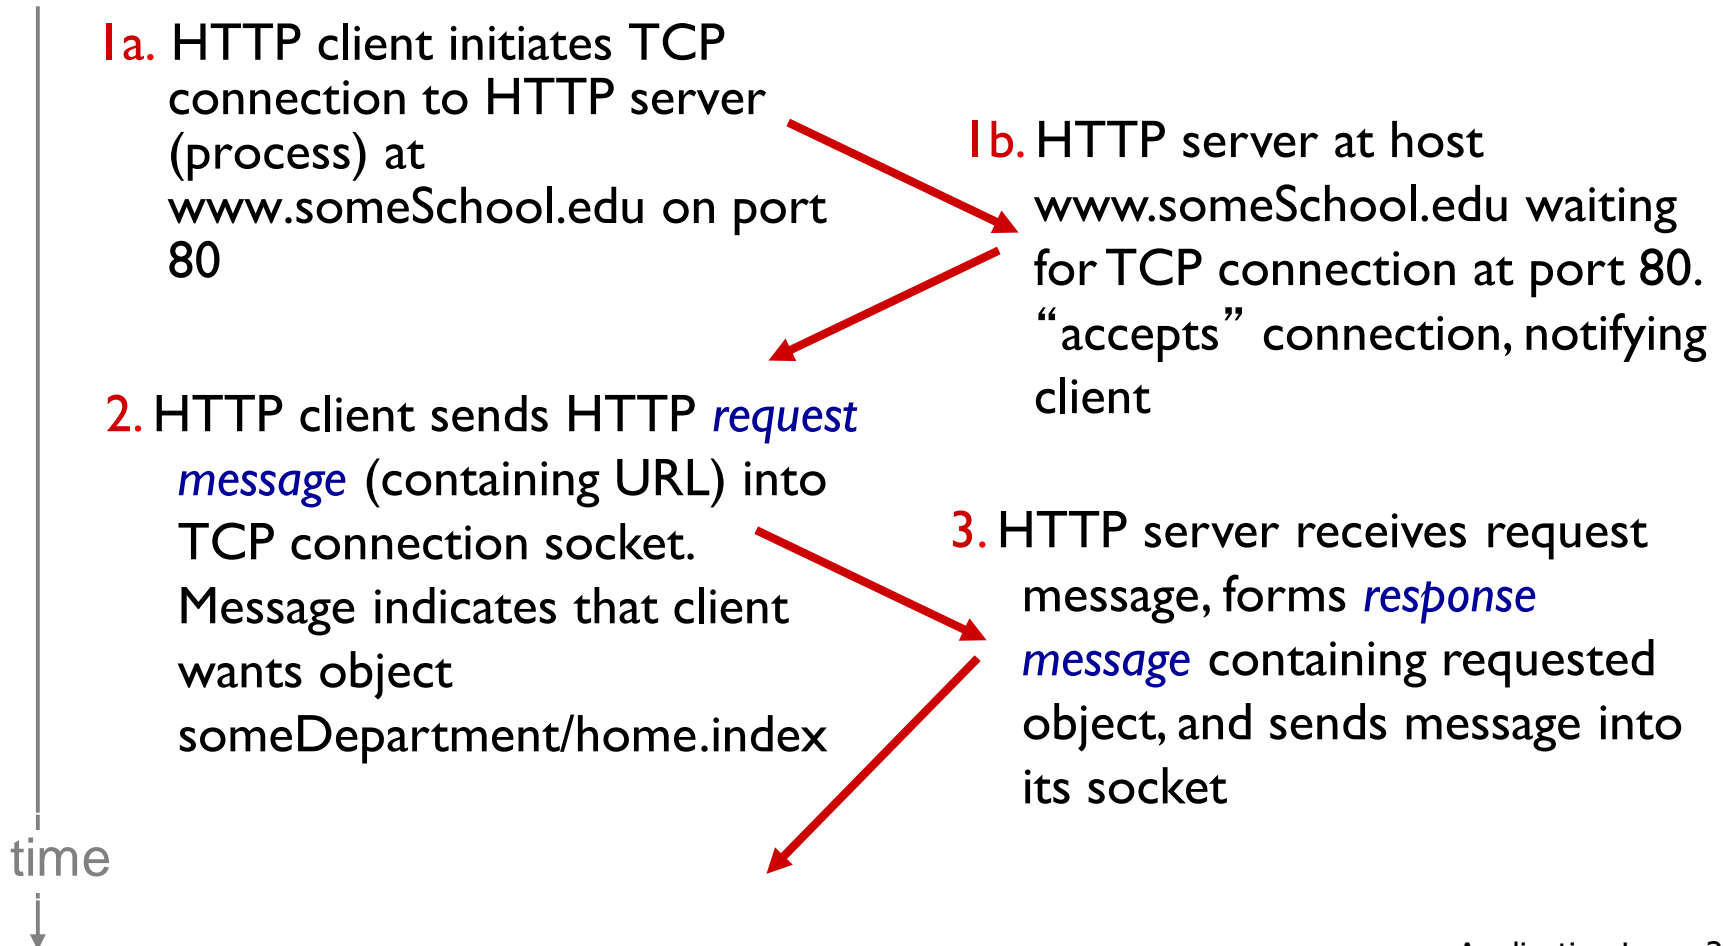
- ❖ multiple objects can be sent over single TCP connection between client, server

# Non-persistent HTTP

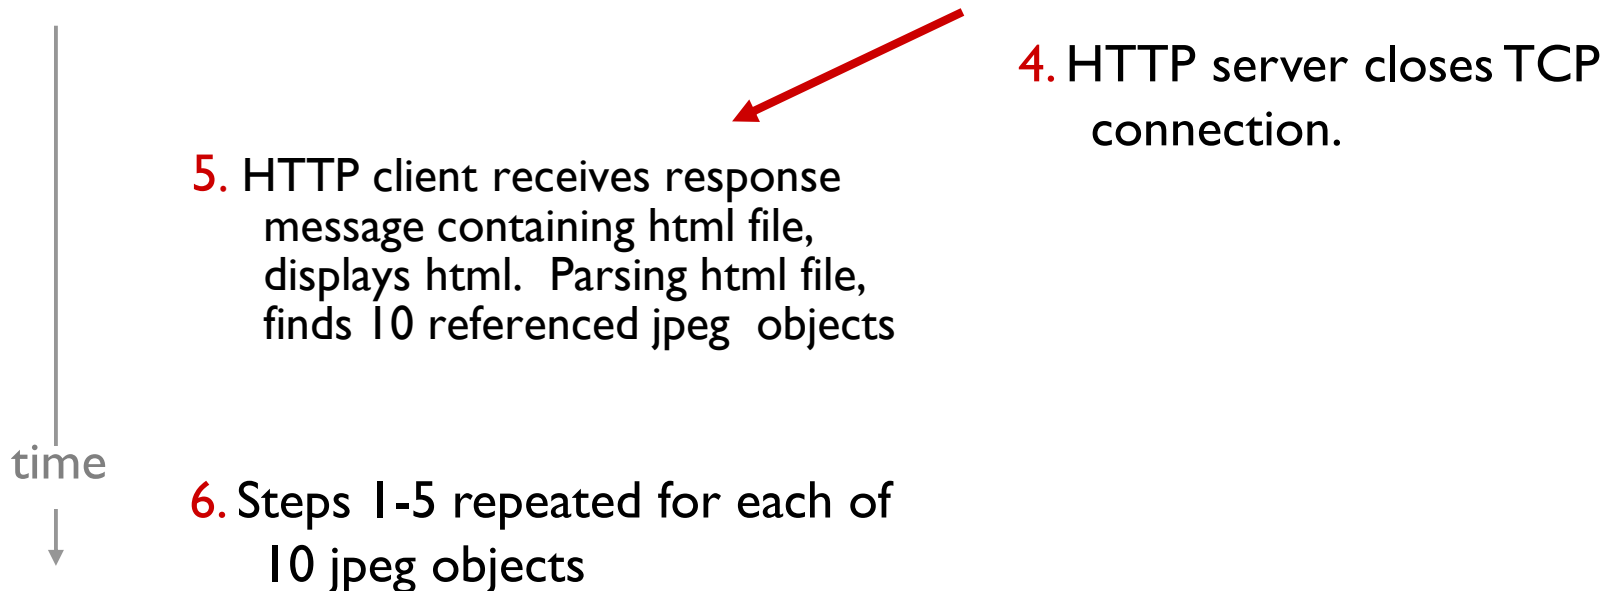
suppose user enters URL:

`www.someSchool.edu/someDepartment/home.index`

(contains text,  
references to 10  
jpeg images)



# Non-persistent HTTP (cont.)




# Non-persistent HTTP: response time

**RTT (definition):** 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 =  
 $2\text{RTT} + \text{file transmission time}$



# 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

## *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 request message

- ❖ two types of HTTP messages: *request, response*
- ❖ **HTTP request message:**
  - ASCII (human-readable format)

request line  
(GET, POST,  
HEAD commands)

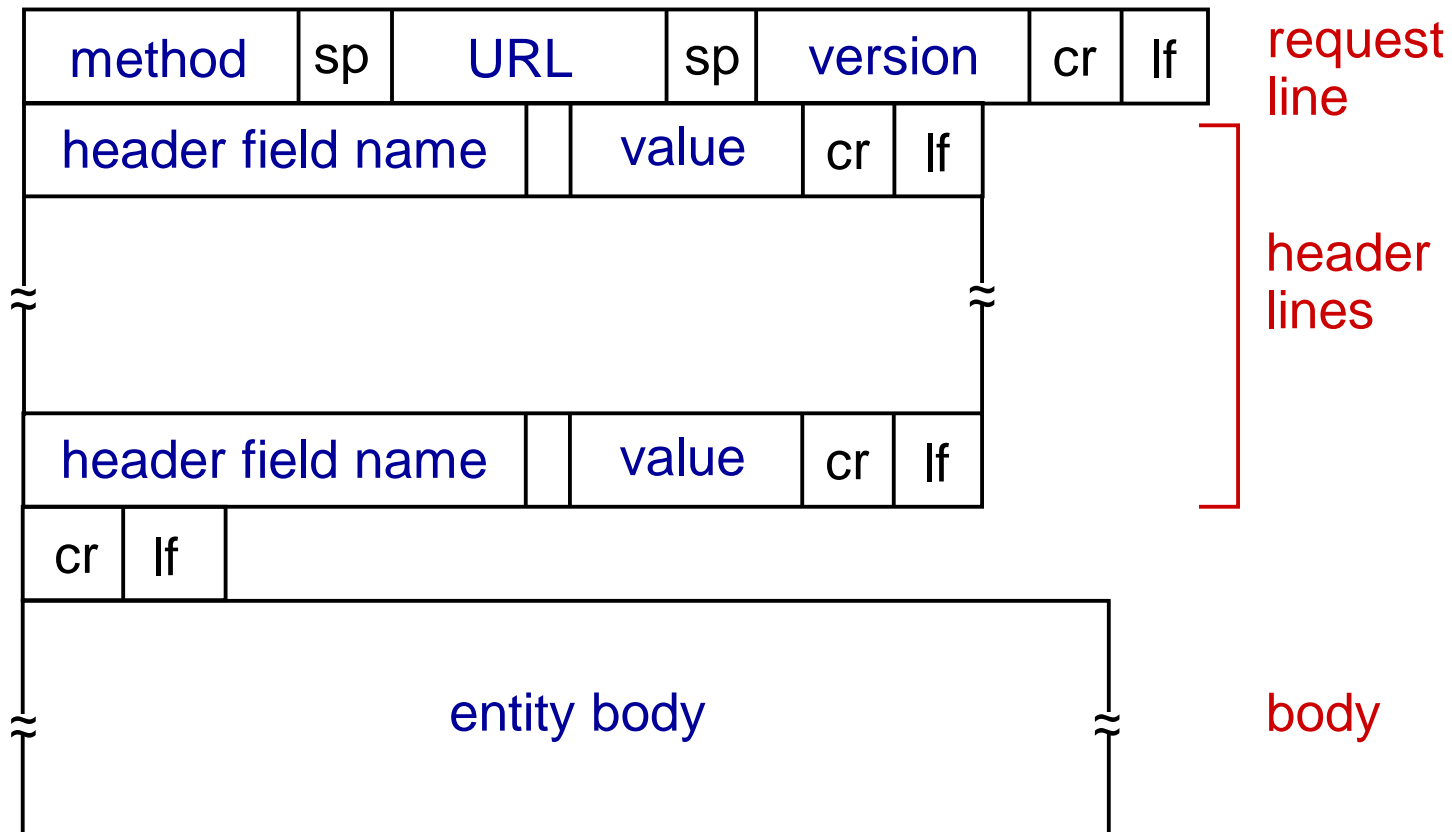
header  
lines

carriage return,  
line feed at start  
of line indicates  
end of header lines

```
GET /index.html HTTP/1.1\r\n
Host: www-net.cs.umass.edu\r\n
User-Agent: Firefox/3.6.10\r\n
Accept: text/html,application/xhtml+xml\r\n
Accept-Language: en-us,en;q=0.5\r\n
Accept-Encoding: gzip,deflate\r\n
Accept-Charset: ISO-8859-1,utf-8;q=0.7\r\n
Keep-Alive: 115\r\n
Connection: keep-alive\r\n
\r\n
```

carriage return character  
line-feed character

# 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`



# 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

# HTTP response message

status line

(protocol

status code

status phrase)

HTTP/1.1 200 OK\r\n

Date: Sun, 26 Sep 2010 20:09:20 GMT\r\n

Server: Apache/2.0.52 (CentOS)\r\n

Last-Modified: Tue, 30 Oct 2007 17:00:02  
GMT\r\n

header  
lines

ETag: "17dc6-a5c-bf716880"\r\n

Accept-Ranges: bytes\r\n

Content-Length: 2652\r\n

Keep-Alive: timeout=10, max=100\r\n

Connection: Keep-Alive\r\n

Content-Type: text/html; charset=ISO-8859-  
1\r\n

\r\n

data, e.g.,  
requested  
HTML file

data data data data data ...

# 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**

# 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)

# 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

# Cookies: keeping "state" (cont.)

client



server



cookie file

usual http request msg

Amazon server  
creates ID  
1678 for user

usual http response  
**set-cookie: 1678**

create  
entry


backend  
database



usual http request msg  
**cookie: 1678**

cookie-  
specific  
action

access



usual http response msg

access

cookie-  
specific  
action

one week later:



usual http request msg  
**cookie: 1678**

usual http response msg

# 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

# 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





# 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)


# 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:*

- ❖ access link utilization = **99%** *problem!*
- ❖ total delay = Internet delay + access delay + LAN delay  
= 2 sec + minutes + usecs




# 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:

- ❖ access link utilization = ~~99%~~ → 9.9%
- ❖ total delay = Internet delay + access delay + LAN delay  
= 2 sec + ~~minutes~~ → msec



**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:*

- ❖ access link utilization = ?
- ❖ total delay = ?

*How to compute link utilization, delay?*


*Cost:* web cache (cheap!)



# 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
  - =  $0.6 * (\text{delay from origin servers}) + 0.4 * (\text{delay when satisfied at cache})$
  - =  $0.6 (2.01) + 0.4 (\sim \text{msecs})$
  - =  $\sim 1.2 \text{ secs}$
  - less than with 154 Mbps link (and cheaper too!)



# Conditional GET

- ❖ **Goal:** don't send object if cache has up-to-date cached version

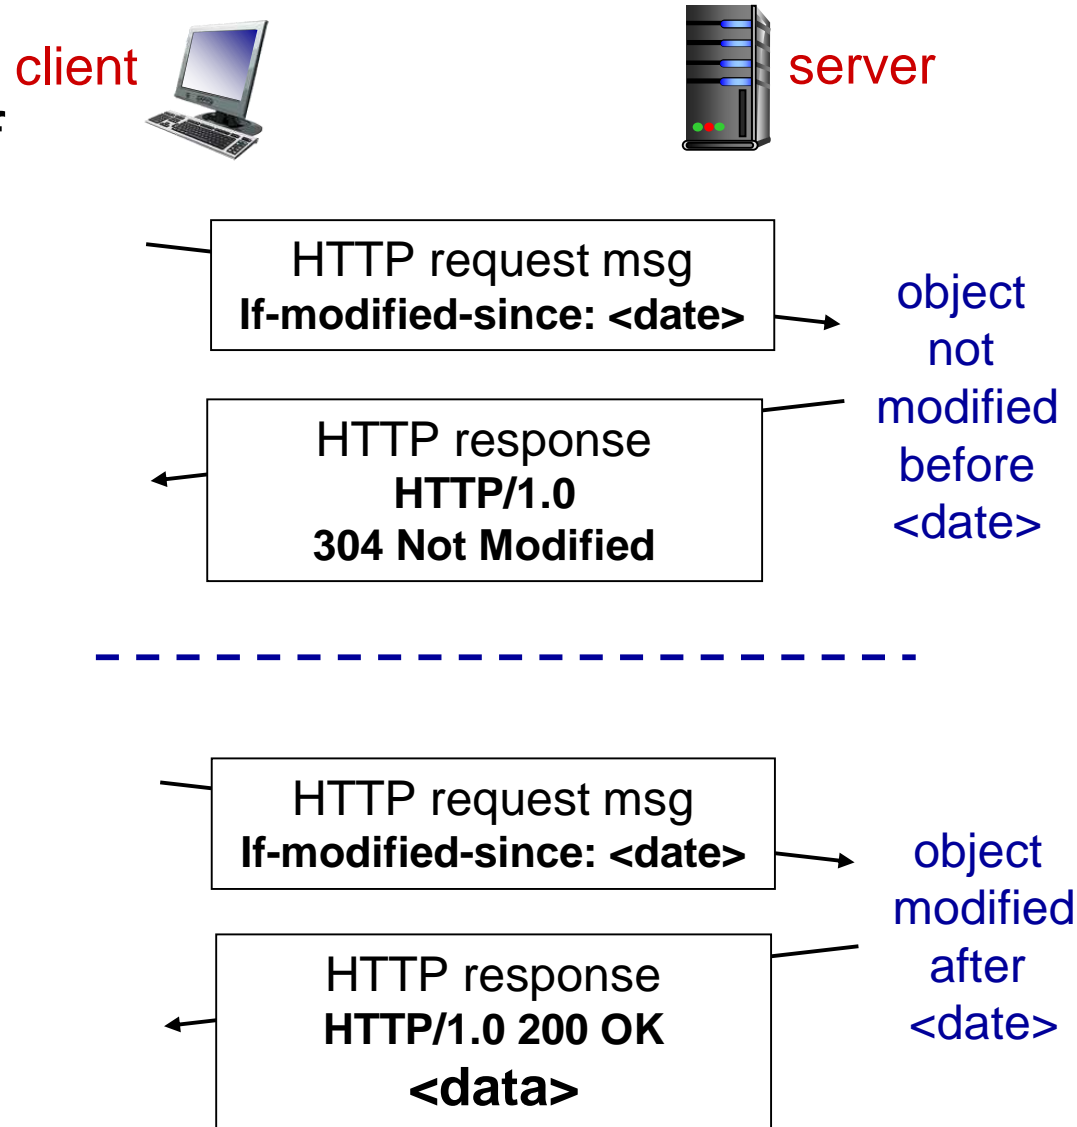
- no object transmission delay
- lower link utilization

- ❖ **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`



# Chapter 2: outline

## 2.1 principles of network applications

- app architectures
- app requirements

## 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 and TCP

# Web Caching and HTTPS

- ❖ Caching as a technique to reduce user (perceived) response time
- ❖ *Who caches?*
  - Origin server (database – memory)
  - Gateway – reverse proxy (shared cache)
  - Proxy (e.g., ISP – share cache)
  - Browser (local to user)
- ❖ *The S stands to TLS, successor of SSL*
  - HTTP over TLS / HTTP over SSL
  - Application security layer guaranteeing privacy, integrity of communication between entities involved in communication
  - Application data are encrypted, middle-boxes can check who is involved in communication but cannot read the data



# Potential Impact of HTTPS

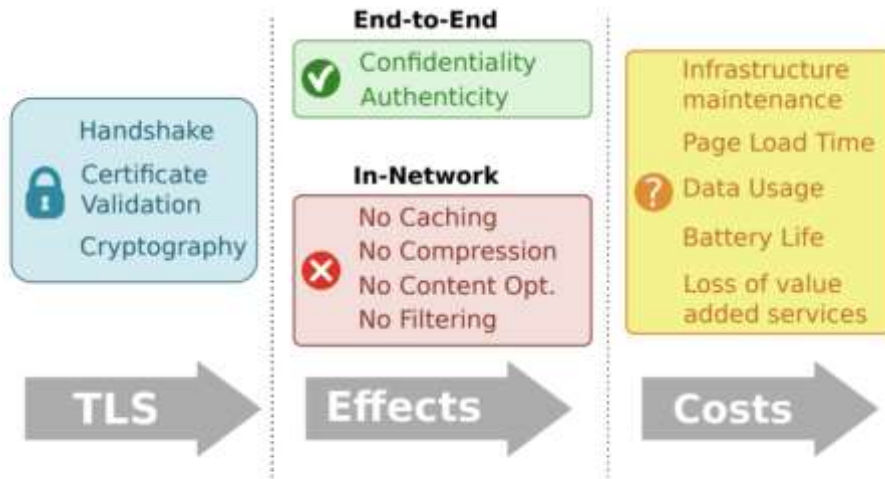


Figure 1: The HTTPS adoption impact chain.

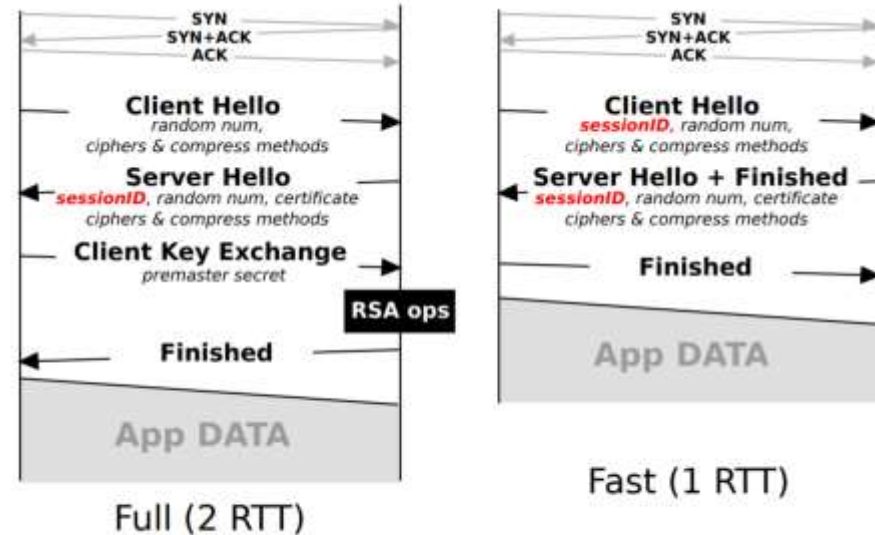


Figure 2: SSL/TLS negotiation.

- ❖ Middle-boxes on the delivery chain cannot act on data anymore – data is encrypted, transparent to the middle-boxes
- ❖ Before worrying, we first need to check how much S in HTTPs is out there?

# How much S in HTTPS

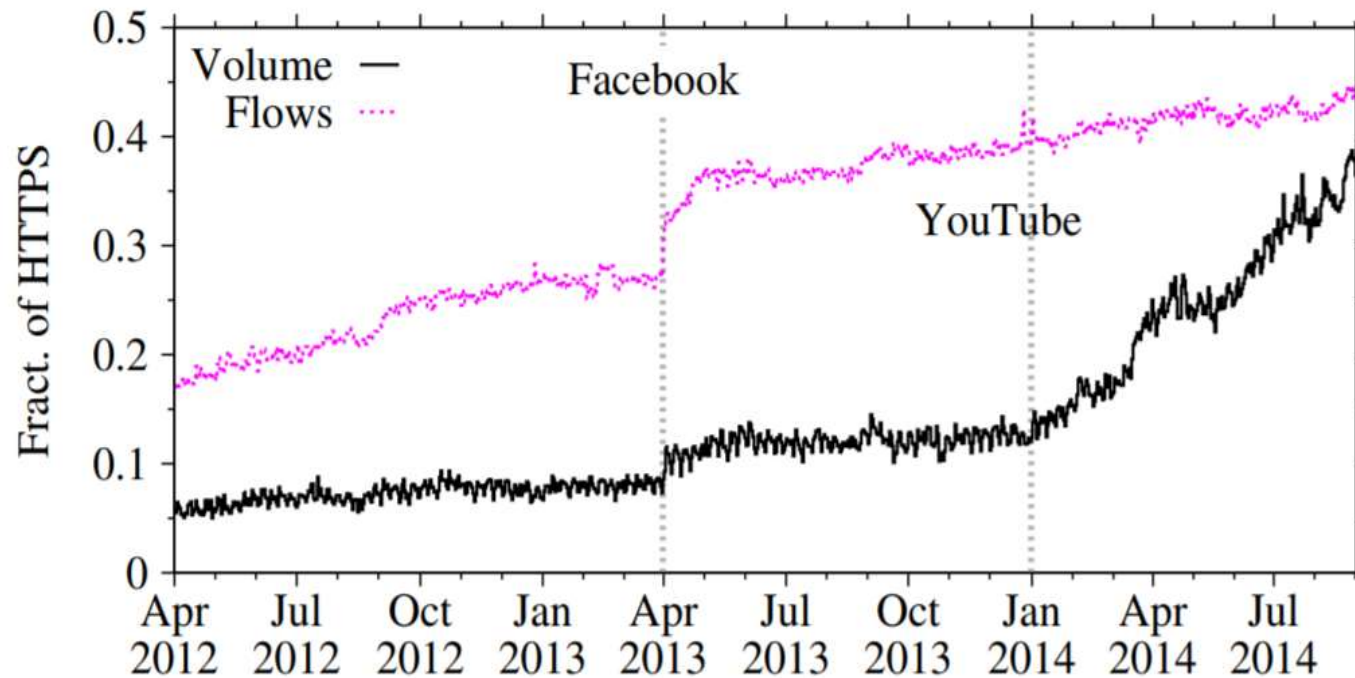


Figure 3. Evolution of HTTPS volume and flow shares over 2.5 years. Vertical lines show the transition to HTTPS for Facebook and YouTube

# How much S in HTTPS




Figure 4. Comparing HTTPS shares over three one-week periods in the Res-ISP dataset. Percentages in the bars highlight year-to-year growth.

# Quantifying the Impact of S

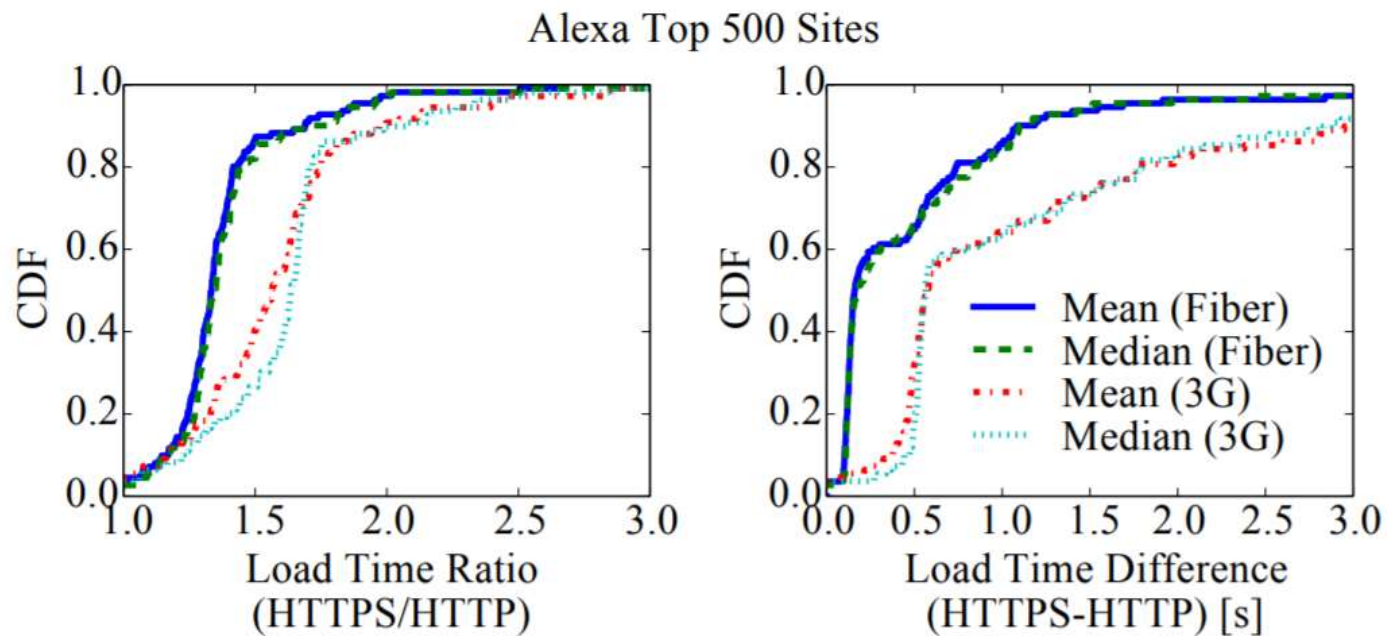


Figure 5: Webpage load time inflation for the Alexa top 500.

# Quantifying the Impact of S

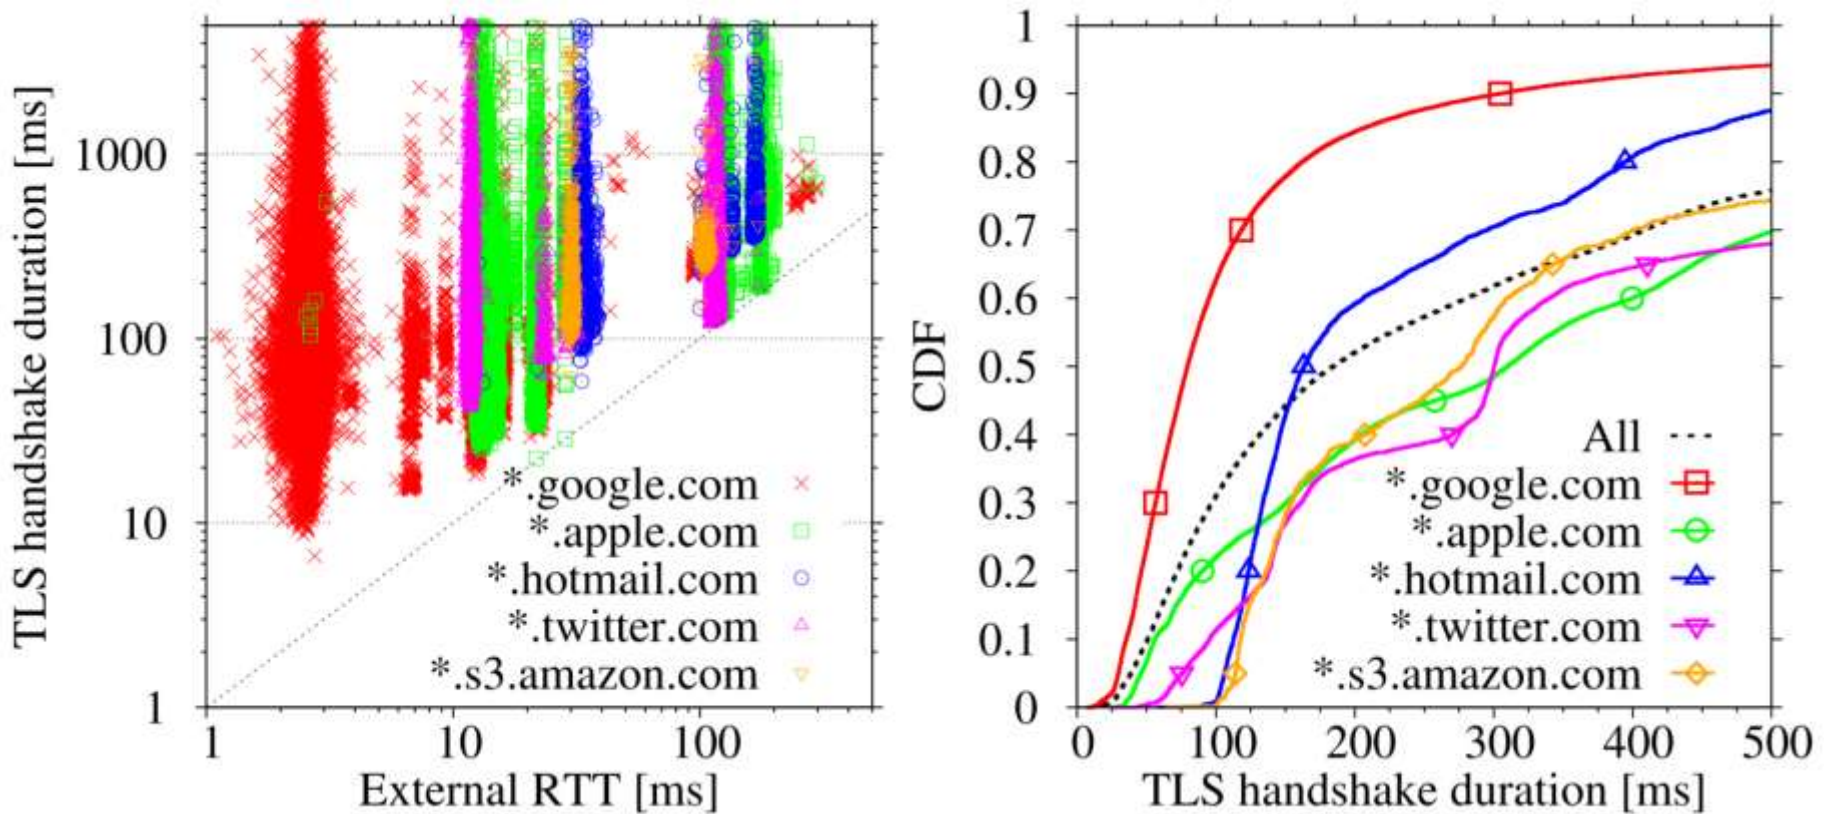


Figure 7: Quantifying TLS handshake costs. Scatter plot of the TLS handshake duration with respect to server distance (left). TLS handshake duration CDF (right).

External RTT: TCP SYN – ACK response

1 Million distinct TLS flows present in the dataset

# Quantifying the Impact of S

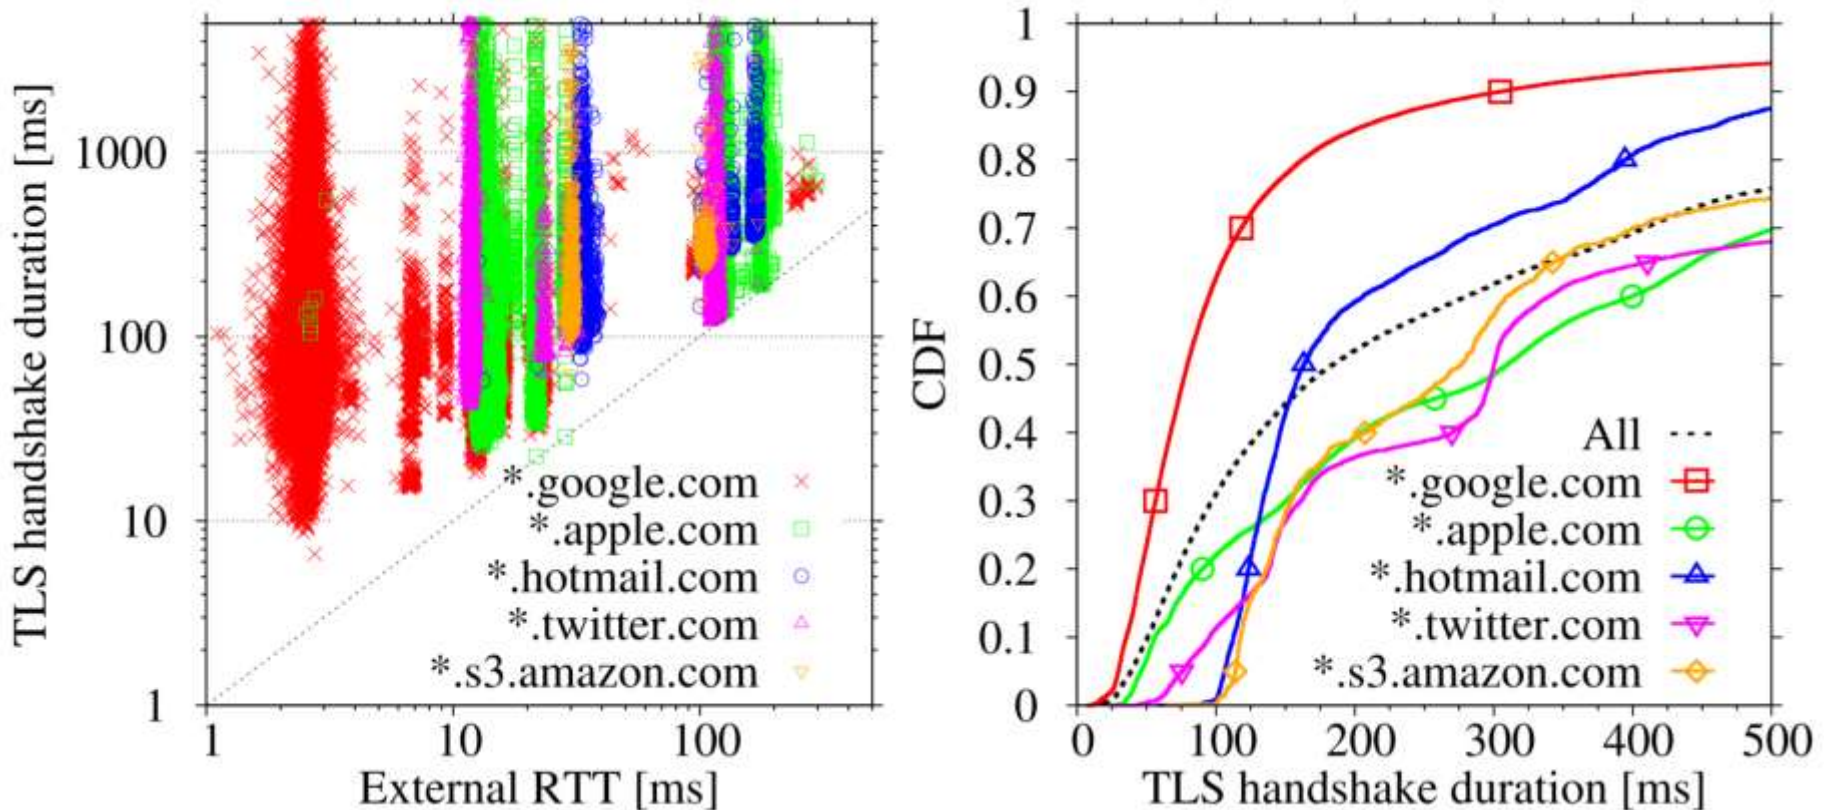


Figure 7: Quantifying TLS handshake costs. Scatter plot of the TLS handshake duration with respect to server distance (left). TLS handshake duration CDF (right).

External RTT: TCP SYN – ACK response

1 Million distinct TLS flows present in the dataset

# Considerations

---

- ❖ HTTPS accounts for 50% of all HTTP connections and is no longer used solely for small objects, suggesting that the cost of deployment is justifiable and manageable for many services.
- ❖ The extra latency introduced by HTTPS is not negligible, especially in a world where one second could cost 1.6 billion in sales (Amazon case study)
- ❖ Most users are unlikely to notice significant jumps in data usage due to loss of compression, but ISPs stand to see a large increase in upstream traffic due to loss of caching
- ❖ From a user perspective browser-caching techniques and proactive content push to the edges are feasible
- ❖ From an operator (ISP, content provider) viewpoint this poses a burden
  - Sol#1: Secure only essential parts of the communication others through HTTP (privacy?)
  - Sol#2: Split proxy approach for not sensitive information (man-in-the-middle (Deep Packet Inspection)?)
  - This said: Youtube mobile version is on clear HTTP