



Computer Networks

Lisa Frye, Instructor

frye@kutztown.edu

Kutztown University

HTTP




Figure 2.6 ♦ HTTP request-response behavior



HTTP Resources / Objects

- File
- Query Result
- Document
- Output of CGI Script
- Etc.

HTTP Characteristics

- Client/Server
- Uses TCP as Transport-Layer Protocol
- Bi-Directional Transfer
- Stateless Protocol
- Capability Negotiation
- Caching

Nonpersistent Connections

- Client sends request (TCP connection created)
- Server sends response
- TCP connection closed

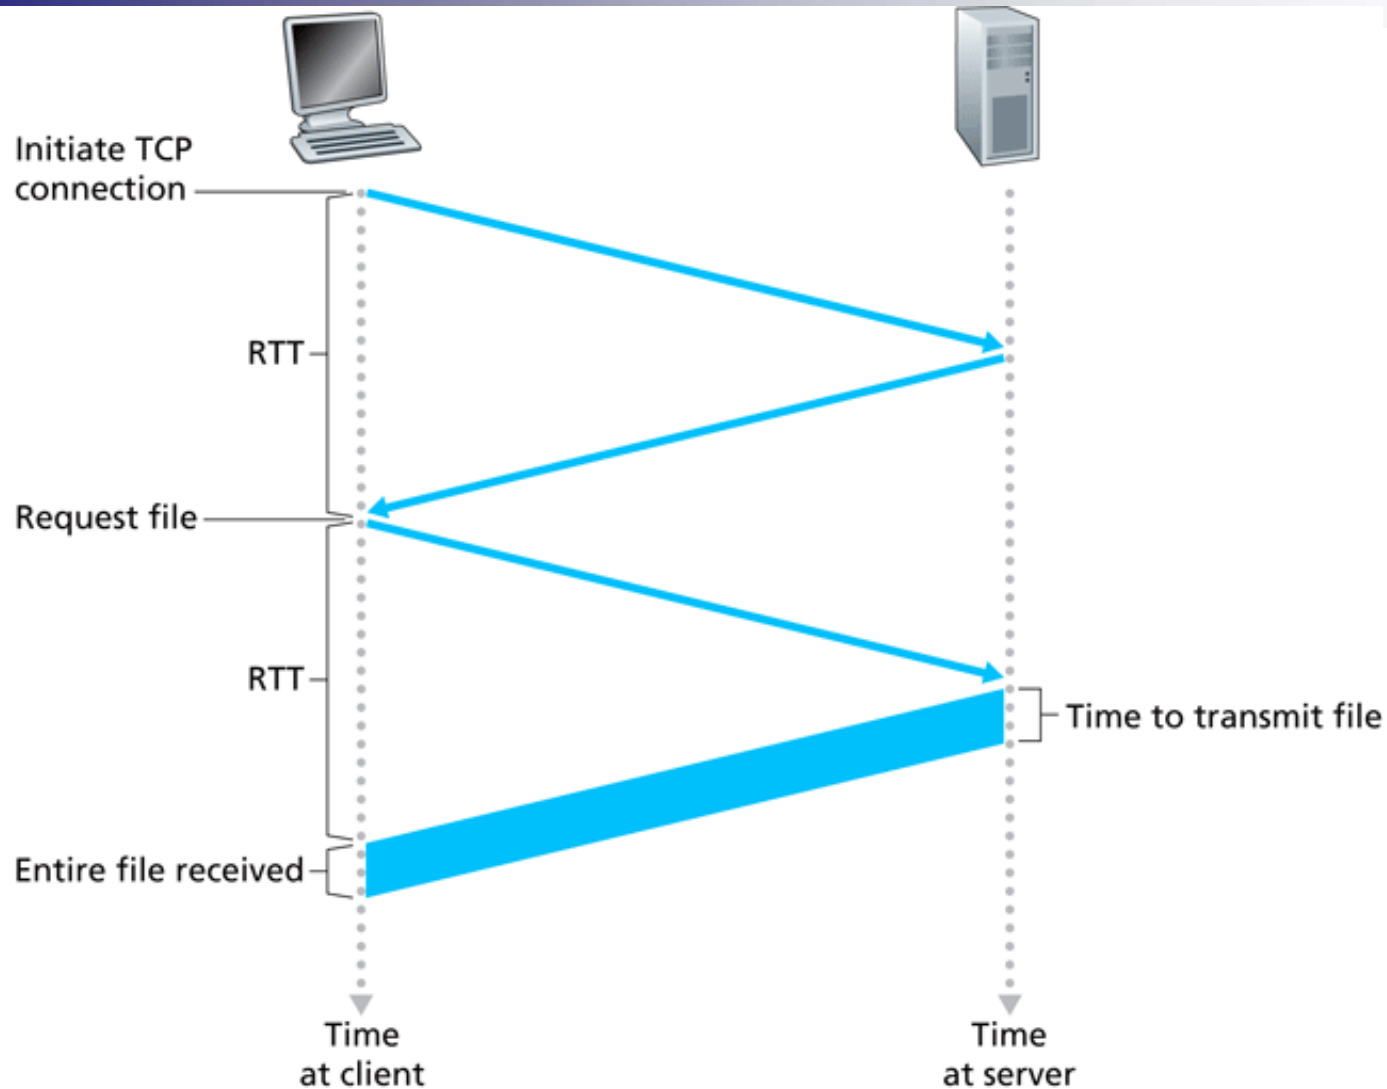


Figure 2.7 ♦ Back-of-the-envelope calculation for the time needed to request and receive an HTML file

Persistent Connections

- TCP connection remains open
 - 1 less RTT
- Pipelining (default HTTP 1.1)
- Without Pipelining

- When a user requests a Web page that consists of some text and two images, the client will send one request message and receive three response messages.
 - True or False?



- With nonpersistent connections between browser and origin server, it is possible for a single TCP segment to carry two distinct HTTP request messages.
 - True or False?



- Two distinct Web pages (for example, www.kutztown.edu/research.html and www.kutztown.edu/students.html) can be sent over the same persistent connection.
 - True or False?



HTTP Request Message

GET /somedir/page.html HTTP/1.1

Host: www.someschool.edu

Connection: Close

User-agent: Mozilla/4.0

Accept-language: en

HTTP Response Message

HTTP/1.1 200 OK

Connection: close

Date: Thu, 07 Jul 2007 12:00:15 GMT

Server: Apache/1.3.0 (Unix)

Last-Modified: Sun, 6 May 2007 09:23:24 GMT

Content-Length: 6821

Content-Type: text/html

(data data data)

HTTP Message Formats

- An initial line.
- Zero or more header lines,
- A blank line (i.e. a CRLF by itself), and
- An optional message body (e.g. a file, or query data, or query output).



Initial Request Line

- Method name.
- Local path.
- Version.

Initial Response Line

- HTTP version.
- Response status code.
 - 1xx indicates an informational message only
 - 2xx indicates success of some kind
 - 3xx redirects the client to another URL
 - 4xx indicates an error on the client's part
 - 5xx indicates an error on the server's part
- Reason phrase.

Header Line

- Provide information about request or response.
- One line per header.
- Same format used for emails.
- Request headers.
- Response headers.

Some Headers

- Content type
- Content length
- Content-Encoding
- Content-Language
- Connection
- Accept-Language
- User-Agent
- Host
- Date
- Server

Sample HTTP Exchange

To retrieve the file at the

URL:`http://www.somehost.com/path/file.html`

1. Open a socket to the host `www.somehost.com`, port 80 (use the default port of 80 because none is specified in the URL).
2. Send something like the following through the socket:

```
GET /path/file.html HTTP/1.0
```

```
From: someuser@jmarshall.com
```

```
User-Agent: HTTPTool/1.0
```

```
[blank line here]
```

The server should respond with something like the following, sent back through the same socket:

HTTP/1.0 200 OK

Date: Fri, 31 Dec 1999 23:59:59 GMT

Content-Type: text/html

Content-Length: 1354

<html>

<body>

<h1>Hello There!</h1>

(more file contents) </body> </html>

Manual HTTP

- From a Unix prompt, open a connection to an HTTP server with something like:
 - `telnet www.somehost.com 80`
- Then enter your request line by line, like:
 - `GET /path/file.html HTTP/1.1`
 - [headers here, if any]
 - [blank line here]
- After you finish your request with the blank line, you'll see the raw response from the server, including the status line, headers, and message body.

Other HTTP Methods

- HEAD

- Return response headers only

- POST

- Send data to server to be processed

- Conditional GET

- If-Modified-Since:



Other HTTP Mechanisms

- Authorization
- Cookies

- Caching



Slow Content Retrieval

- Slow link
- Congested link
- Slow server

Content Distribution

- Replicating content on multiple servers in the Internet
- Providing requesting end systems a means to determine the servers that can deliver the content the fastest



Content Distribution Schemes

- Web caching
- Content distribution networks (CDNs)
- Peer-to-peer file sharing

Proxy Servers




Figure 2.11 ♦ Clients requesting objects through a Web cache

Proxy Server Example




Figure 2.12 ♦ Bottleneck between an institutional network and the Internet

Example

- Internet connection = 15 Mbps
- Average object size = 1 Mb
- Average request rate = 15 requests/second
- “Internet Delay” = 2 seconds

Example, cont.

- Total response time = LAN delay + access delay + Internet delay
- Traffic intensity on LAN = $(15 \text{ requests/sec}) * (1000 \text{ kbits/request}) / (100\text{Mbps}) = 0.15 \rightarrow \text{negligible}$
- Traffic intensity on access link = $(15 \text{ requests/sec}) * (1000 \text{ kbits/request}) / (15\text{Mbps}) = 1$

Solution #1

- Faster access link
 - 100 Mbps
 - Total response time = 2 seconds

Solution #2





Figure 2.13 ♦ Adding a cache to the institutional network

- 
- Which solution is less expensive?
 1. Faster access link
 2. Web cache



CDN

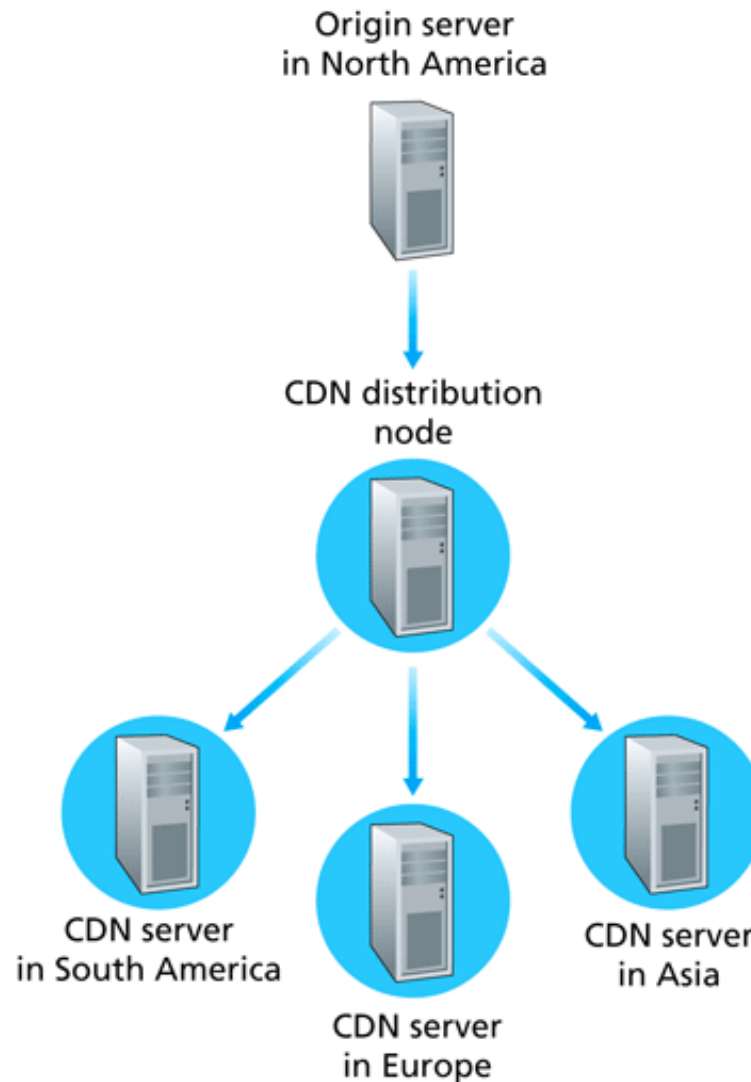


Figure 7.8 ♦ The CDN pushes content provider's tagged objects to its CDN servers.

Finding the “best” server

- DNS redirection
- Internal network map
- Keep track of “best” server for ISPs

FTP




Figure 2.14 ♦ FTP moves files between local and remote file systems

FTP Connections

- Parallel Connection

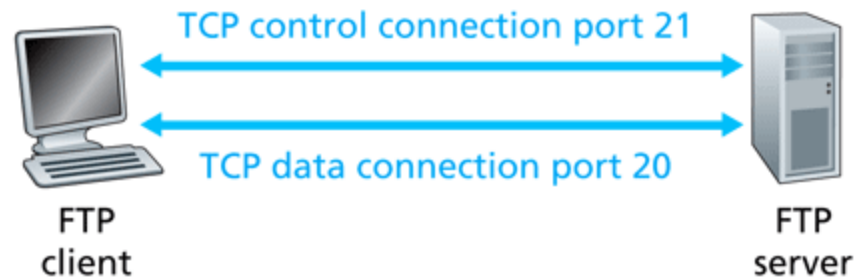



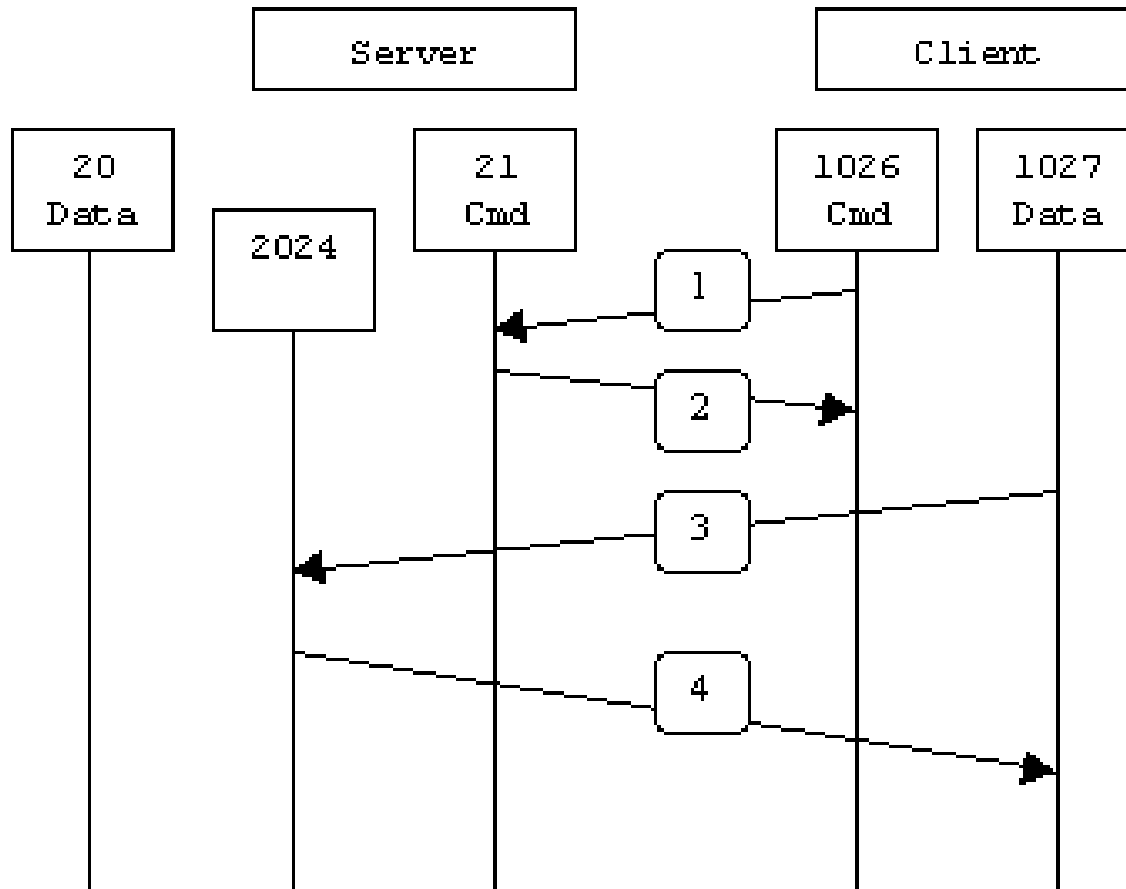
Figure 2.15 ♦ Control and data connections

- Out-of-band (FTP)
- In-band (HTTP)

Active FTP



Passive FTP





FTP State Information

- User account
- User's current directory
- Session constraints

FTP Commands

- USER username
- PASS password
- LIST
- RETR filename
- STOR filename

FTP Replies

- 331 Username OK, password required.
- 125 Data connection already open; transfer starting.
- 425 Can't open data connection.
- 452 Error writing file.