



COMP 4021

Internet Computing

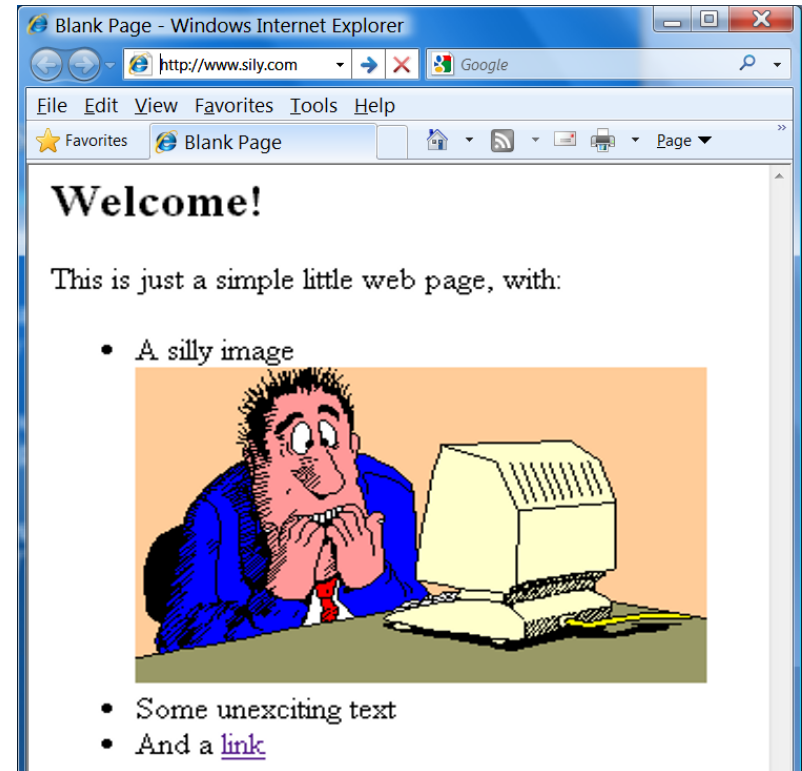
The Browser Process / HTTP

Dr. Kenneth LEUNG

Slides created by Dr. David ROSSITER

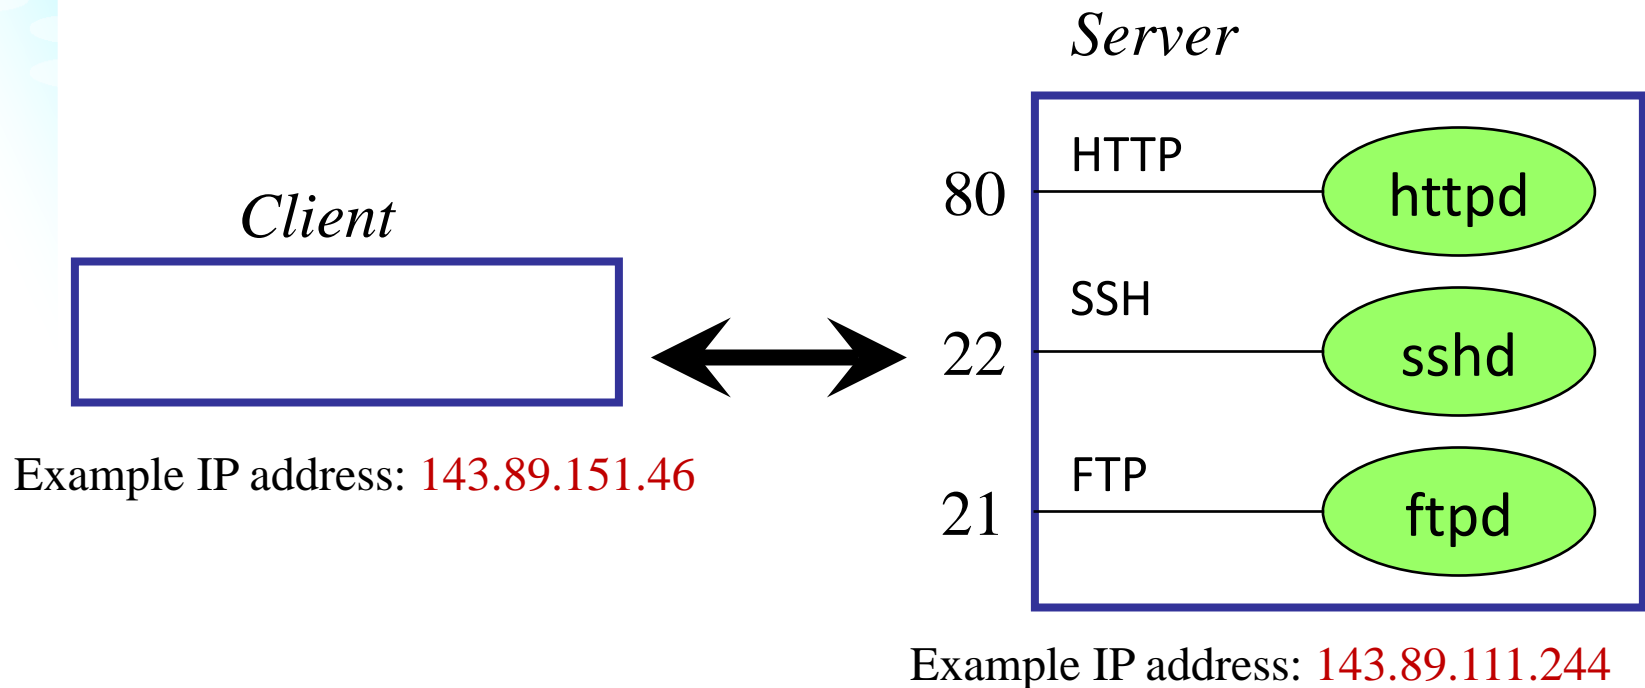
A Browser Uses HTTP

- Here is a simple web page
- Assume this web page is at `http://www.silly.com`
- To get the page, you type the URL in the browser and press Enter
- The browser requests the web page using *HTTP*



IP Address and Ports

- May processes may be running on a server
- They each use a different *port* (=door)



Example Ports

<i>Protocol</i>	<i>Default Port</i>
FTP	21
Telnet	23
HTTP	80
NNTP (Usenet news)	119
ICQ	5190
Quake game	26000
Half-life game	27010

Client-Server Communication

- The browser connects to the machine silly.com using the HTTP protocol
- No port was specified by the user so the browser assumes port 80



Client



HTTP Protocol



Server

Client's Request



- The message (called a *request*) that the browser sends to the silly.com server at port 80 is:

GET / HTTP/1.1

You could simply send this line only.

User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1)

Could send this also, meaning that:
The program is IE, running on Windows XP

```
HTTP/1.1 200 OK
Date: Mon, 1 Nov 2010 09:09:47 HKT
Server: Apache/1.3.6 (Unix) mod_ssl/2.2.8 OpenSSL/0.9.2b
Last-Modified: Mon, 14 Apr 2008 09:39:08 HKT
Accept-Ranges: bytes
Content-Length: 507
Keep-Alive: timeout=15, max=100
Connection: Keep-Alive
Content-Type: text/html
```

```
<!doctype html public "-//w3c//dtd html 4.0 transitional//en">
<html>
<head>
  <title>Simple page</title>
</head>
<body>
<h2>Welcome!</h2>
This is just a simple little web page, with:
<ul>
<li>
A silly image 
</li>
<li>Some unexciting text</li>
<li>
And a <a href="http://www.winniethepooh.com">link</a></li>
</ul>
</body>
</html>
```

Response header

Server's Response

The silly.com
server responds
with everything
shown here

The file follows
the header



Server's Response - Header



HTTP/1.1 200 OK

This line tells the client what version of the HTTP protocol the server uses and says that the document has been found and is going to be transmitted.

Date: Mon, 1 Nov 2010 09:09:47 HKT

Current date on the server in Greenwich Mean Time (GMT)

Server: Apache/1.3.6 (Unix) mod_ssl/2.2.8 OpenSSL/0.9.2b

Tells the client what type of software the server is running, in this case Apache, version 1.3.6, running under Unix

Last-Modified: Mon, 14 Apr 2008 09:39:08 HKT

Tells the client the last time that the document was modified

Content-Length: 507

Tells the client how many bytes are coming

Content-Type: text/html

Tells the client the type of the document

Client's Request



. . .

A silly image ``

. . .

The browser sees this in the HTML, and understands that the web page also needs an image file. So it sends a request for the image:

```
GET /man.gif HTTP/1.1
```

```
User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1)
```

Server's Response



HTTP/1.1 200 OK

Date: Mon, 1 Nov 2010 09:09:48 HKT

Server: Apache/1.3.6 (Unix) mod_ssl/2.2.8 OpenSSL/0.9.2b

Last-Modified: Mon, 14 Apr 2008 09:39:12 HKT

Content-Length: 4627

Content-Type: image/gif

The browser knows that GIF data follows

v({CCPP P]]]kkkxxx;;;®®®»»»ÉÉÉÖÖÖääääñññÿÿÿÿÿÿÿS?U1¶D@âU
nQÿÿ{.(/{Rn;/,ÿÿ H° *\È Ì F(±;Å3jÜÈ±£Ç 5VD(âÃG
. . .



This is the GIF data (it looks strange because it is not meant for viewing as text)

Client

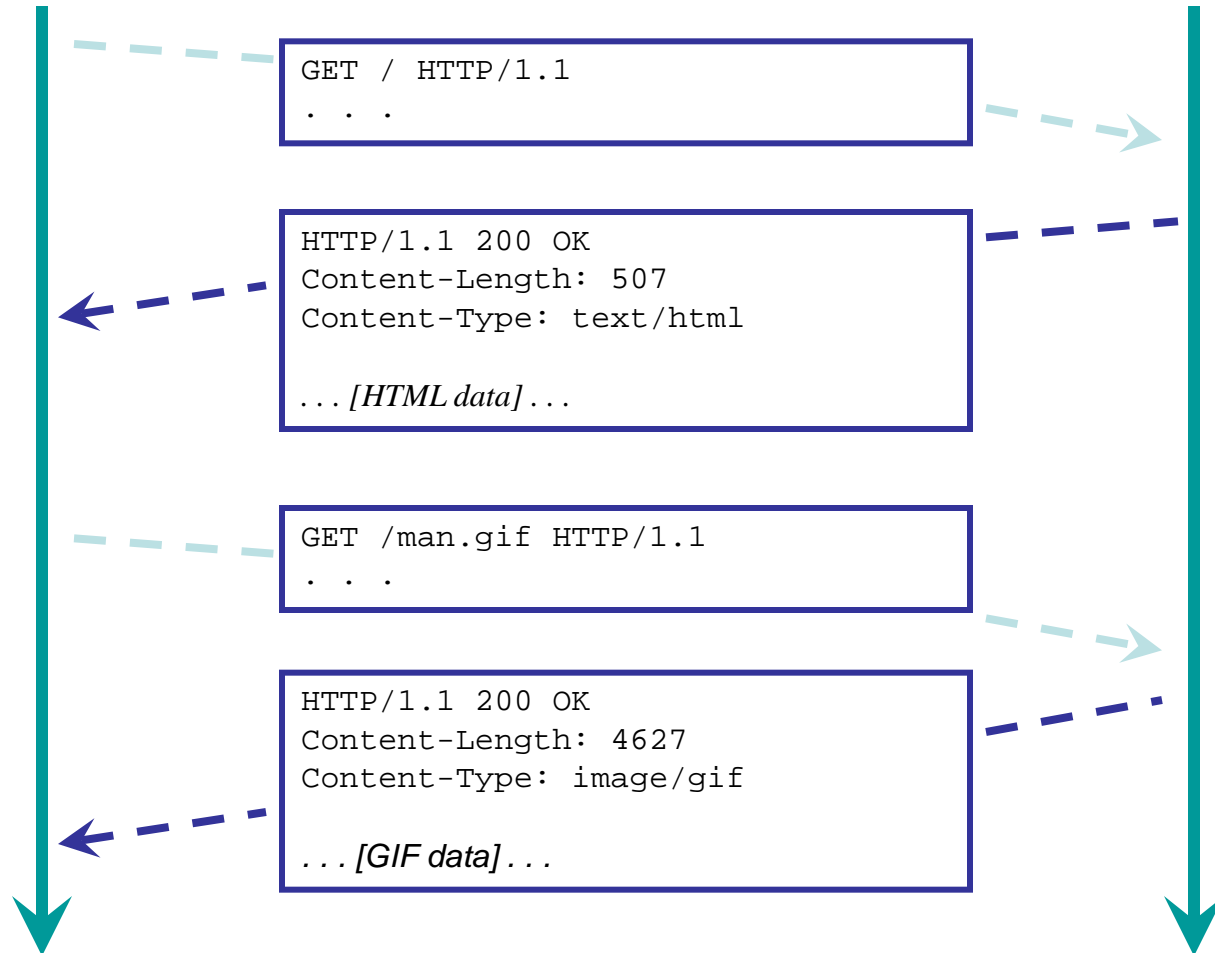


Summary

Server



Time



Forms



- When you submit a form the browser sends the form data to the server, as well as the name of the program on the server which it needs to be given to



Movie Search

Select the name and/or the year of the movie you want to search for.

Title:

Year:

Press submit when you're ready

HTML Source Code

```
<html>

<head>
<title>Movie Database!</title>
</head>

<form method="post"
action="http://ihome.ust.hk/~rossiter/cgi-bin/show_environment.php">

<h1>Movie Search</h1>
Select the name and/or the year of the movie you want to search
  for.<p />
Title:<input type="text" name="movie_title" value="" />
Year:<input type="text" name="movie_year" value="" />

<br />
<br />
Press submit when you're ready

<input type="submit" value="Submit">

</form>
</html>
```

Sending Form Data

- After the *Submit* button is pressed, the browser connects to the server shown in the 'action' field, using port 80

`action="http://ihome.ust.hk/~rossiter/cgi-bin/show_environment.php"`

- In this case the server is ihome.ust.hk
- The browser then sends:

`POST /~rossiter/cgi-bin/show_environment.php`

`User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1)`

`Content-type: application/x-www-form-urlencoded`

`Content-length: 35`

35 bytes to follow.

`movie_title=spiderman+3&movie_year=`

Two parameter pairs separated by a '&', with spaces replaced by '+'

The browser now includes this line, specifying the type of data being sent. Previously this wasn't necessary, as no content was sent with the request.



Response from Server

- After receiving the client request, the client will give all the sent information to the server side program `show_environment.php`
- The program does whatever it is programmed to do
- Probably, it outputs something back to the client - whatever it outputs (i.e. prints) goes straight back to the browser