



Distributed Computing

Lesson 10: HTTP

Thomas Weise · 汤卫思

tweise@hfu.edu.cn · <http://www.it-weise.de>

Hefei University, South Campus 2
Faculty of Computer Science and Technology
Institute of Applied Optimization
230601 Shushan District, Hefei, Anhui, China
Econ. & Tech. Devel. Zone, Jinxiu Dadao 99

合肥学院 南艳湖校区/南2区
计算机科学与技术系
应用优化研究所
中国 安徽省 合肥市 蜀山区 230601
经济技术开发区 锦绣大道99号

1 Browser and Web Server

2 HTTP




website


- We now know HTML, web pages, and how they can be “located” in the internet via URLs.
- But how does the content of a web page come from the web server to our computer?
- We will learn about HTTP, the protocol existing for this purpose.
- How is HTTP related to TCP sockets and what we’ve learned so far?

- **Client/Server** systems are the most common application structure in the internet and corporate networks


- **Client/Server** systems are the most common application structure in the internet and corporate networks




- **Client/Server** systems are the most common application structure in the internet and corporate networks





- **Client/Server** systems are the most common application structure in the internet and corporate networks

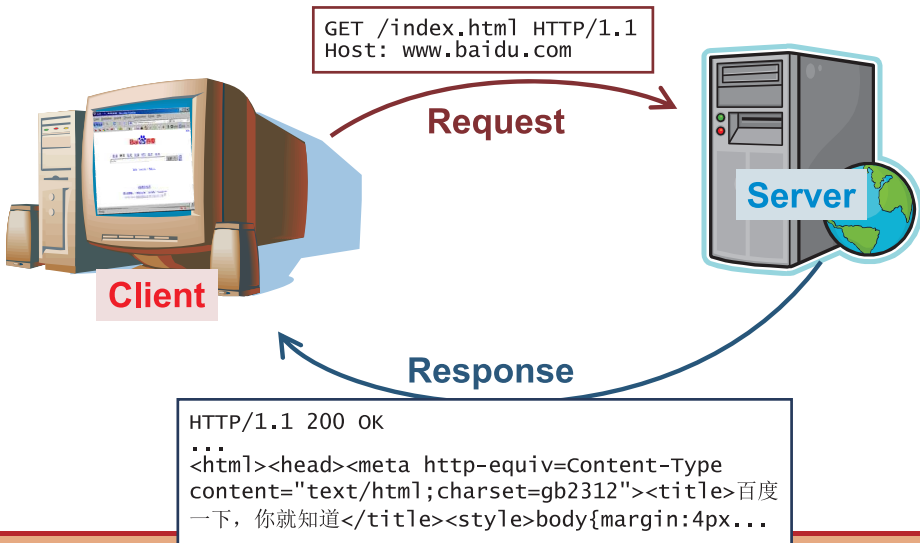


- **Client/Server** systems are the most common application structure in the internet and corporate networks











The screenshot shows the official website of the University of Science and Technology of China (USTC). The browser window is titled "中国科学技术大学" and the address bar shows "www.ustc.edu.cn". The website features a large blue globe graphic with the text "网络课堂" (Network Classroom), "报考科大" (Apply to USTC), "科大校友" (USTC Alumni), and "在校师生" (Students and Teachers). Below the globe, there is a navigation menu with links such as "学校概况" (School Overview), "院系介绍" (Department Introduction), "师资队伍" (Faculty), "本科生教育" (Undergraduate Education), "研究生教育" (Graduate Education), "科大文化" (USTC Culture), "科学研究" (Scientific Research), "发展规划" (Development Planning), "人才招聘" (Talent Recruitment), "校园门户" (Campus Portal), "公共服务" (Public Service), and "电子邮件" (Email).

The main content area is divided into two columns. The left column is titled "科大要闻" (USTC News) and lists several news items with dates. The right column is titled "公告通知" (Announcements) and lists various notices with dates. At the bottom of the page, there is a footer with the text "Copyright 2006-2011 中国科学技术大学 All Rights Reserved" and "联系我们" (Contact Us).

Main Resource:
HTML Code (hypertext)
index.html



Main Resource:
HTML Code (hypertext)
index.html

Resource:
embedded image
all.jpg



Main Resource:
HTML Code (hypertext)
index.html

Resource:
embedded image
all.jpg

7 Resources:
embedded images
img_1.jpg...img_7.jpg



The screenshot shows the USTC website with the following elements:

- Header:** "中国科学技术大学" (University of Science and Technology of China) and "www.ustc.edu.cn".
- Navigation:** Links for "网络课堂", "报考科大", "科大校友", and "在校生".
- Main Content:** A large globe graphic with the text "中国科学技术大学" and "University of Science and Technology of China". Below it are links for "学校概况", "院系介绍", "师资队伍", "本科生教育", "研究生教育", "科大文化", "科学研究", "发展规划", "人才招聘", "校园门户", "公共服务", and "电子邮箱".
- News Section:** A list of news items with dates, including "学校召开2013年校务工作会议", "学校召开全校教授干部大会", "传达全国“两会”精神", "侯建国:加强协同创新 实现人才质量和科研能力双提升", "自然出版集团发布《2012年东亚地区自然出版指数》", "中国科大蝉联中国高校第一", "李卫平教授出席中国工程院院士论坛", "学校举办第七届范旭东奖学金颁奖典礼", "中国教育报中国科大:为学生讲学多开几扇窗", and "我校2013年硕士研究生招生考试工作全面展开".
- Announcements:** A list of announcements with dates, including "一周会议表", "关于评选中国科学技术大学“全院办校、所系结合”工作先进个人和优秀个人的通知", "关于中国科学院2013年度中联合培养博士生项目“选课工作的通知", "关于教职工在两处以上取得个人收入尽快办理纳税申报的通知", "“流年碎影——庆祝中国科大建校50周年”微博主题征文活动启事", "关于2013年清明节放假和师生员工外出春游有关问题的通知", "电信中心西区主干电缆断接通知", and "关于申报崔英东教育基金会2013年高等院校青年教师基金及青年教师科研启动基金的通知".
- Footer:** Copyright 2006-2011 中国科学技术大学 All Rights Reserved. 联系我们 0550-2528号.

Main Resource:
HTML Code (hypertext)
index.html

Resource:
embedded image
all.jpg

Resource:
CSS Style Sheet file
main.css

7 Resources:
embedded images
img_1.jpg...img_7.jpg



Main Resource:
HTML Code (hypertext)
index.html


Resource:
embedded image
all.jpg


Resource:
CSS Style Sheet file
main.css

Resource:
JavaScript file
menu.js

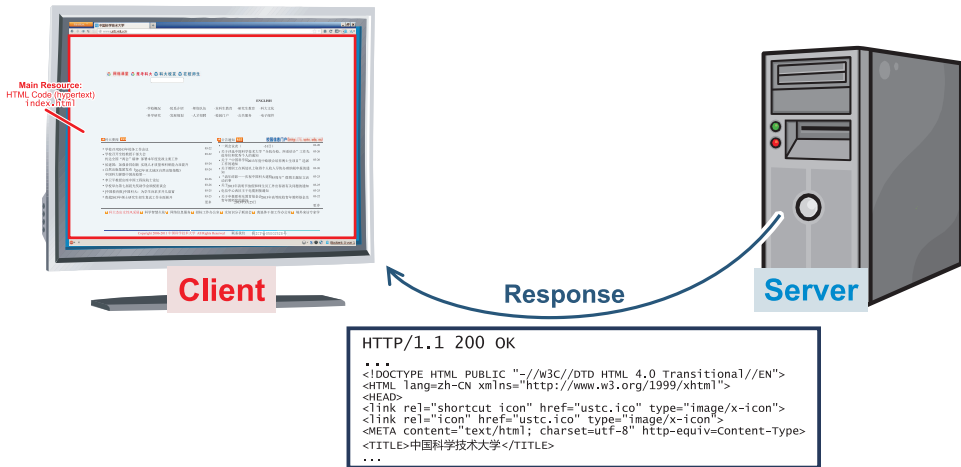
7 Resources:
embedded images
img_1.jpg...img_7.jpg







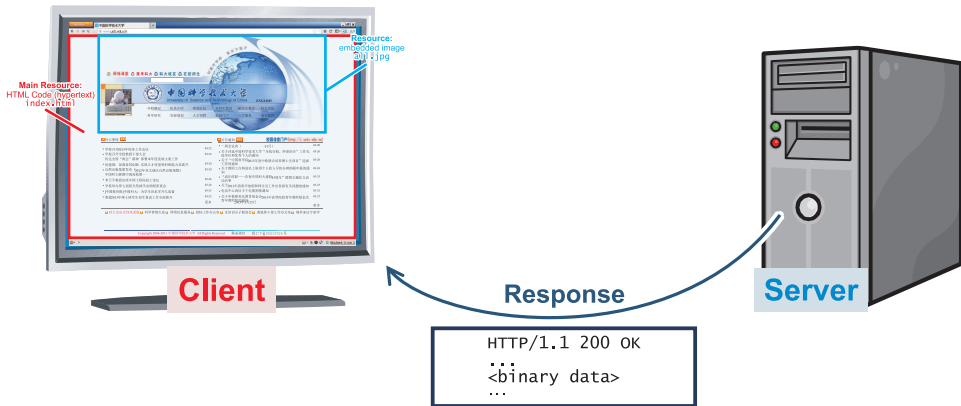
Browser loads Website



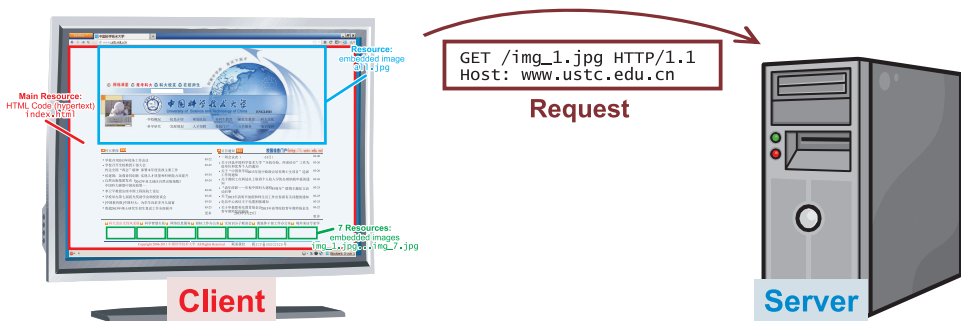
Browser loads Website



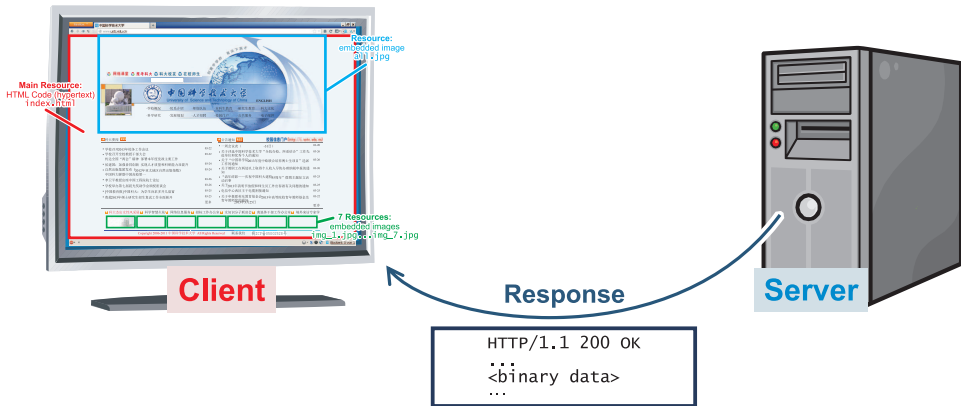
Browser loads Website



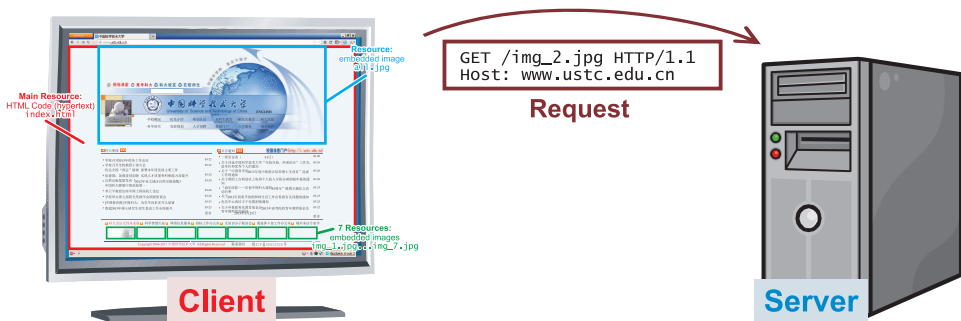
Browser loads Website



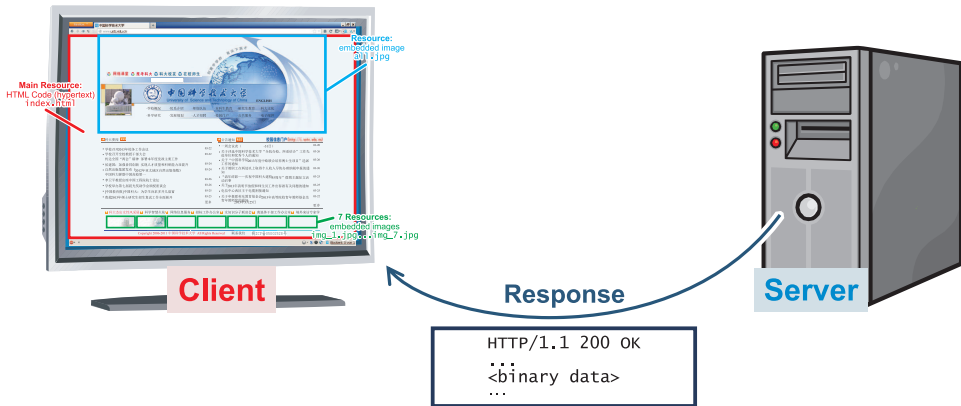
Browser loads Website



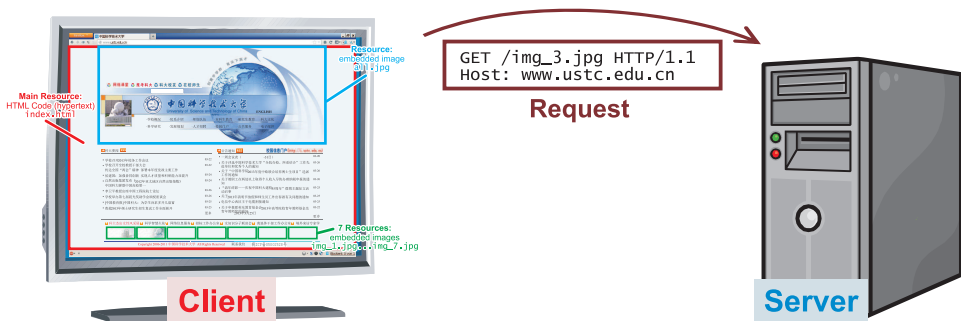
Browser loads Website



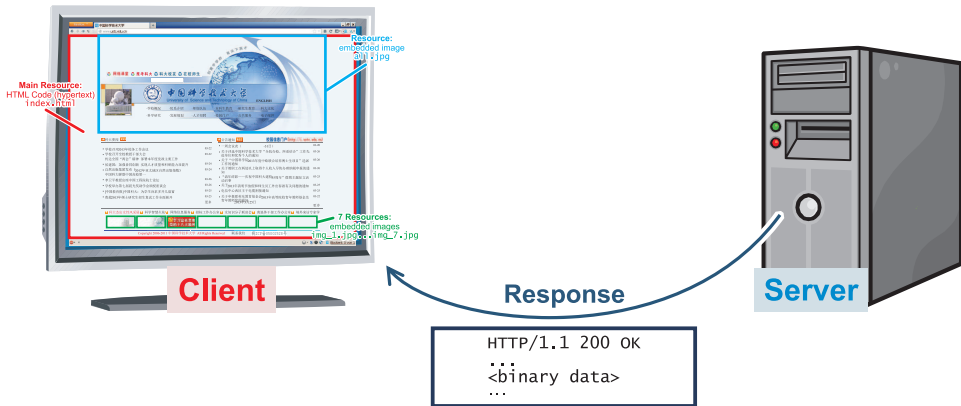
Browser loads Website




Browser loads Website



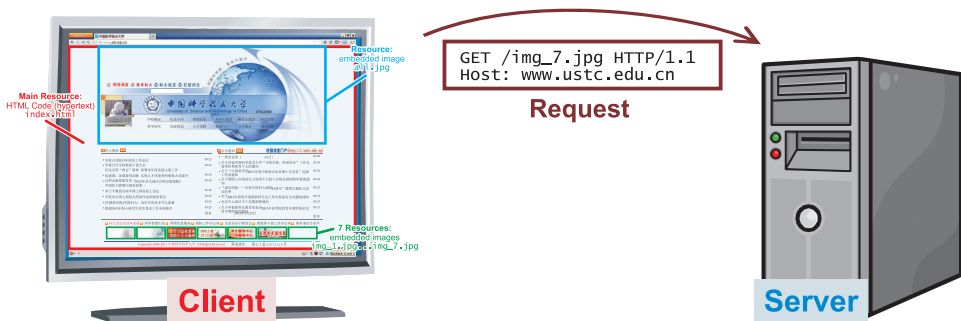
Browser loads Website



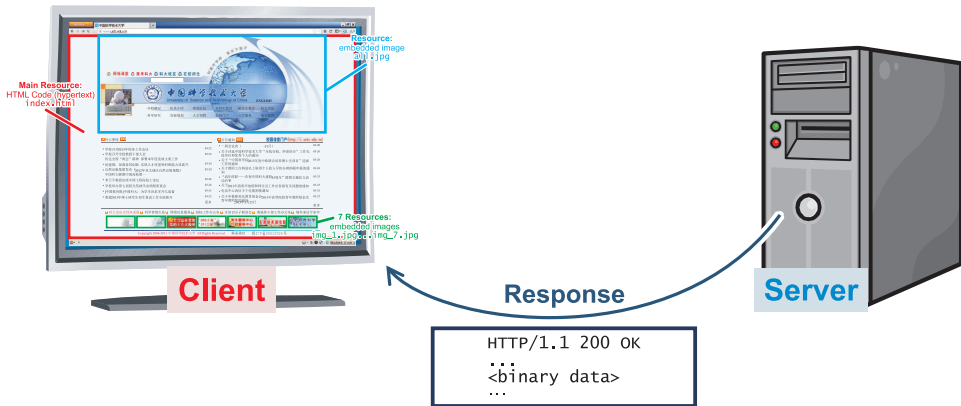
Browser loads Website



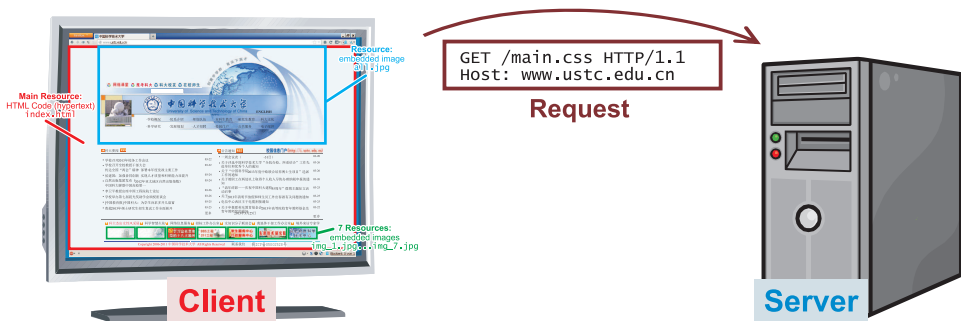
Browser loads Website



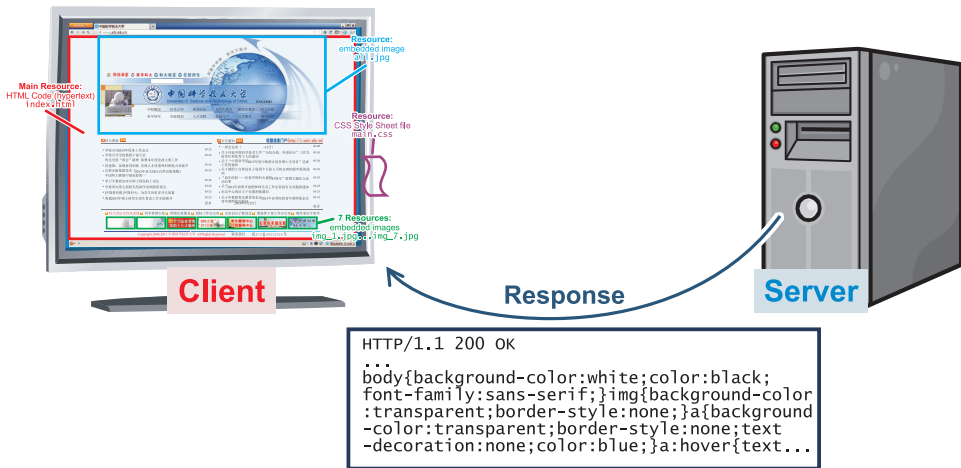
Browser loads Website



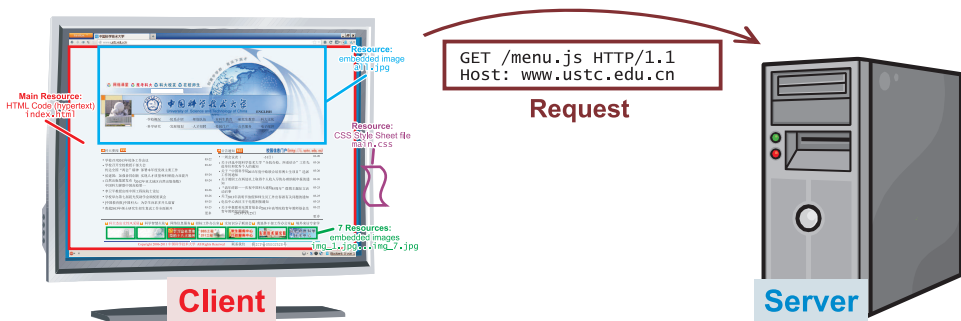
Browser loads Website



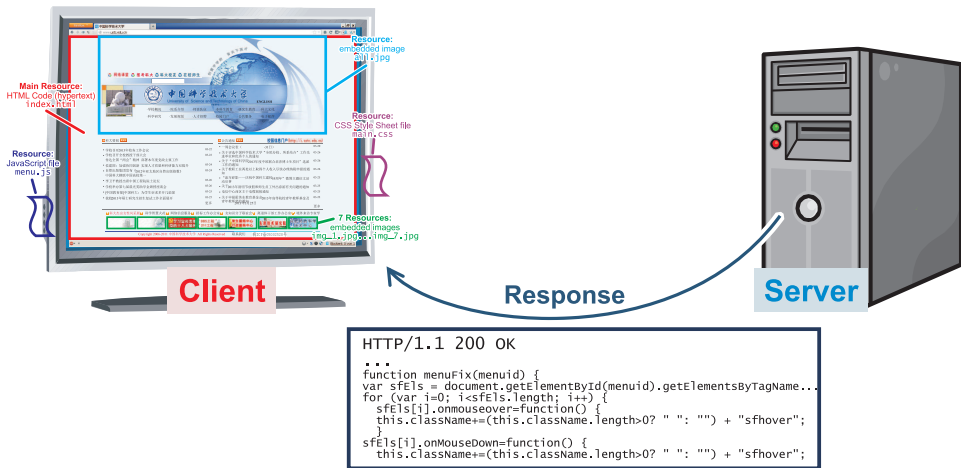
Browser loads Website



Browser loads Website



Browser loads Website



- But how do web browser and web server communicate?

- But how do web browser and web server communicate?
- Well, we already know the Socket API

- But how do web browser and web server communicate?
- Well, we already know the Socket API which gives us access to transport layer protocols such as TCP and UDP

- But how do web browser and web server communicate?
- Well, we already know the Socket API which gives us access to transport layer protocols such as TCP and UDP
- Using TCP seems to be a good idea

- But how do web browser and web server communicate?
- Well, we already know the Socket API which gives us access to transport layer protocols such as TCP and UDP
- Using TCP seems to be a good idea: We transfer files, and files can be characterized as streams

- But how do web browser and web server communicate?
- Well, we already know the Socket API which gives us access to transport layer protocols such as TCP and UDP
- Using TCP seems to be a good idea: We transfer files, and files can be characterized as streams
- But TCP is not enough


- But how do web browser and web server communicate?
- Well, we already know the Socket API which gives us access to transport layer protocols such as TCP and UDP
- Using TCP seems to be a good idea: We transfer files, and files can be characterized as streams
- But TCP is not enough: The web browser must, at least, also send the path/object part of the URL

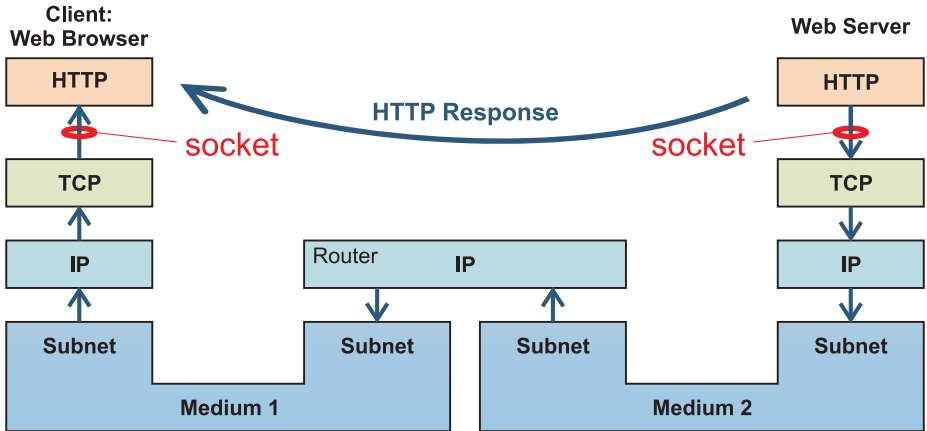
- But how do web browser and web server communicate?
- Well, we already know the Socket API which gives us access to transport layer protocols such as TCP and UDP
- Using TCP seems to be a good idea: We transfer files, and files can be characterized as streams
- But TCP is not enough: The web browser must, at least, also send the path/object part of the URL, so that the web server knows *which* page or ressource the browser wants. . .

- But how do web browser and web server communicate?
- Well, we already know the Socket API which gives us access to transport layer protocols such as TCP and UDP
- Using TCP seems to be a good idea: We transfer files, and files can be characterized as streams
- But TCP is not enough: The web browser must, at least, also send the path/object part of the URL, so that the web server knows *which* page or ressource the browser wants. . . . and the server must be able to send errors or other answers

- But how do web browser and web server communicate?
- Well, we already know the Socket API which gives us access to transport layer protocols such as TCP and UDP
- Using TCP seems to be a good idea: We transfer files, and files can be characterized as streams
- But TCP is not enough: The web browser must, at least, also send the path/object part of the URL, so that the web server knows *which* page or ressource the browser wants. and the server must be able to send errors or other answers
- Protocol on top of TCP, designed to achieve this

- But how do web browser and web server communicate?
- Well, we already know the Socket API which gives us access to transport layer protocols such as TCP and UDP
- Using TCP seems to be a good idea: We transfer files, and files can be characterized as streams
- But TCP is not enough: The web browser must, at least, also send the path/object part of the URL, so that the web server knows *which* page or resource the browser wants. . . . and the server must be able to send errors or other answers
- Protocol on top of TCP, designed to achieve this: HTTP





- HyperText Transfer Protocol (HTTP)

- HyperText Transfer Protocol (HTTP)
- Application-layer protocol for communication between web browser and web server

- HyperText Transfer Protocol (HTTP)
- Application-layer protocol for communication between web browser and web server
- Uses reliable, bi-directional, ordered byte stream provided by TCP

- HyperText Transfer Protocol (HTTP)
- Application-layer protocol for communication between web browser and web server
- Uses reliable, bi-directional, ordered byte stream provided by TCP
- Features

- HyperText Transfer Protocol (HTTP)
- Application-layer protocol for communication between web browser and web server
- Uses reliable, bi-directional, ordered byte stream provided by TCP
- Features:
 - text-based and human readable

- HyperText Transfer Protocol (HTTP)
- Application-layer protocol for communication between web browser and web server
- Uses reliable, bi-directional, ordered byte stream provided by TCP
- Features:
 - text-based and human readable
 - state-less (no information preserved between different HTTP requests)

- HyperText Transfer Protocol (HTTP)
- Application-layer protocol for communication between web browser and web server
- Uses reliable, bi-directional, ordered byte stream provided by TCP
- Features:
 - text-based and human readable
 - state-less (no information preserved between different HTTP requests)
 - anonymous

- HyperText Transfer Protocol (HTTP)
- Application-layer protocol for communication between web browser and web server
- Uses reliable, bi-directional, ordered byte stream provided by TCP
- Features:
 - text-based and human readable
 - state-less (no information preserved between different HTTP requests)
 - anonymous
- HTTP 1.0 ^[1]

- HyperText Transfer Protocol (HTTP)
- Application-layer protocol for communication between web browser and web server
- Uses reliable, bi-directional, ordered byte stream provided by TCP
- Features:
 - text-based and human readable
 - state-less (no information preserved between different HTTP requests)
 - anonymous
- HTTP 1.0 ^[1]:
 - Messages in the MIME (Multipurpose Internet Mail Extensions) format

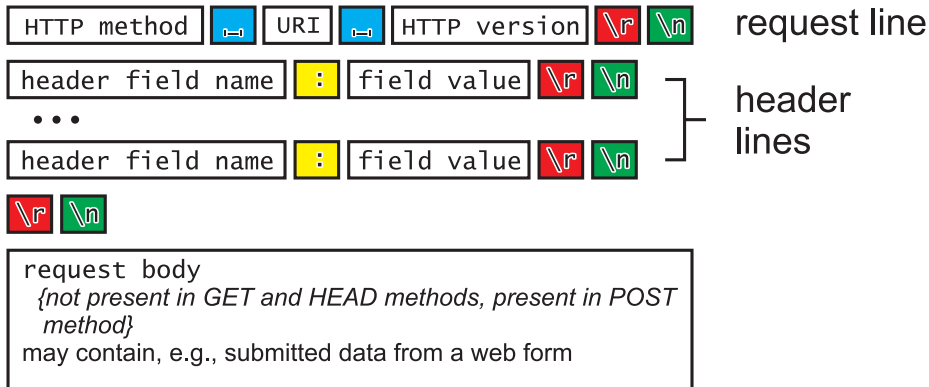
- HyperText Transfer Protocol (HTTP)
- Application-layer protocol for communication between web browser and web server
- Uses reliable, bi-directional, ordered byte stream provided by TCP
- Features:
 - text-based and human readable
 - state-less (no information preserved between different HTTP requests)
 - anonymous
- HTTP 1.0 ^[1]:
 - Messages in the MIME (Multipurpose Internet Mail Extensions) format
- HTTP 1.1 ^[2]

- HyperText Transfer Protocol (HTTP)
- Application-layer protocol for communication between web browser and web server
- Uses reliable, bi-directional, ordered byte stream provided by TCP
- Features:
 - text-based and human readable
 - state-less (no information preserved between different HTTP requests)
 - anonymous
- HTTP 1.0 ^[1]:
 - Messages in the MIME (Multipurpose Internet Mail Extensions) format
- HTTP 1.1 ^[2]:
 - Support for hierarchical proxies, caching, persistent connections, virtual hosts

- HyperText Transfer Protocol (HTTP)
- Application-layer protocol for communication between web browser and web server
- Uses reliable, bi-directional, ordered byte stream provided by TCP
- Features:
 - text-based and human readable
 - state-less (no information preserved between different HTTP requests)
 - anonymous
- HTTP 1.0 ^[1]:
 - Messages in the MIME (Multipurpose Internet Mail Extensions) format
- HTTP 1.1 ^[2]:
 - Support for hierarchical proxies, caching, persistent connections, virtual hosts
 - Determine applications' capabilities

- HyperText Transfer Protocol (HTTP)
- Application-layer protocol for communication between web browser and web server
- Uses reliable, bi-directional, ordered byte stream provided by TCP
- Features:
 - text-based and human readable
 - state-less (no information preserved between different HTTP requests)
 - anonymous
- HTTP 1.0 ^[1]:
 - Messages in the MIME (Multipurpose Internet Mail Extensions) format
- HTTP 1.1 ^[2]:
 - Support for hierarchical proxies, caching, persistent connections, virtual hosts
 - Determine applications' capabilities
 - TCP connections can be re-used

- HyperText Transfer Protocol (HTTP)
- Application-layer protocol for communication between web browser and web server
- Uses reliable, bi-directional, ordered byte stream provided by TCP
- Features:
 - text-based and human readable
 - state-less (no information preserved between different HTTP requests)
 - anonymous
- HTTP 1.0 ^[1]:
 - Messages in the MIME (Multipurpose Internet Mail Extensions) format
- HTTP 1.1 ^[2]:
 - Support for hierarchical proxies, caching, persistent connections, virtual hosts
 - Determine applications' capabilities
 - TCP connections can be re-used (good for inline resources/images in HTML)



HTTP method



URI



HTTP version

\r

\n

- HTTP method

HTTP method



URI



HTTP version



- HTTP method:
 - **GET**: download the data that belongs to a specified resource (can include parameters), e.g., a web page



- HTTP method:
 - **GET** : download the data that belongs to a specified resource (can include parameters), e.g., a web page
 - **HEAD** : returns only the headers of a GET response (no request body)

HTTP method



URI



HTTP version



- HTTP method:

- **GET** : download the data that belongs to a specified resource (can include parameters), e.g., a web page
- **HEAD** : returns only the headers of a GET response (no request body)
- **POST** : data is sent in the request body that should be handed to the requested resource, maybe the data filled into a form or a file



- HTTP method:
 - GET : download the data that belongs to a specified resource (can include parameters), e.g., a web page
 - HEAD : returns only the headers of a GET response (no request body)
 - POST : data is sent in the request body that should be handed to the requested resource, maybe the data filled into a form or a file
 - PUT : upload a representation of the specified resource



- HTTP method:
 - GET : download the data that belongs to a specified resource (can include parameters), e.g., a web page
 - HEAD : returns only the headers of a GET response (no request body)
 - POST : data is sent in the request body that should be handed to the requested resource, maybe the data filled into a form or a file
 - PUT : upload a representation of the specified resource
 - DELETE : delete the specified resource



- HTTP method:
 - **GET** : download the data that belongs to a specified resource (can include parameters), e.g., a web page
 - **HEAD** : returns only the headers of a GET response (no request body)
 - **POST** : data is sent in the request body that should be handed to the requested resource, maybe the data filled into a form or a file
 - **PUT** : upload a representation of the specified resource
 - **DELETE** : delete the specified resource
- relative part of URI



- HTTP method:
 - **GET** : download the data that belongs to a specified resource (can include parameters), e.g., a web page
 - **HEAD** : returns only the headers of a GET response (no request body)
 - **POST** : data is sent in the request body that should be handed to the requested resource, maybe the data filled into a form or a file
 - **PUT** : upload a representation of the specified resource
 - **DELETE** : delete the specified resource
- relative part of URI: (contains, e.g., path/object)



- HTTP method:
 - **GET** : download the data that belongs to a specified resource (can include parameters), e.g., a web page
 - **HEAD** : returns only the headers of a GET response (no request body)
 - **POST** : data is sent in the request body that should be handed to the requested resource, maybe the data filled into a form or a file
 - **PUT** : upload a representation of the specified resource
 - **DELETE** : delete the specified resource
- relative part of URI: (contains, e.g., path/object) and identifies a resource relative to server



- HTTP method:
 - GET : download the data that belongs to a specified resource (can include parameters), e.g., a web page
 - HEAD : returns only the headers of a GET response (no request body)
 - POST : data is sent in the request body that should be handed to the requested resource, maybe the data filled into a form or a file
 - PUT : upload a representation of the specified resource
 - DELETE : delete the specified resource
- relative part of URI: (contains, e.g., path/object) and identifies a resource relative to server
- HTTP version



- HTTP method:
 - GET : download the data that belongs to a specified resource (can include parameters), e.g., a web page
 - HEAD : returns only the headers of a GET response (no request body)
 - POST : data is sent in the request body that should be handed to the requested resource, maybe the data filled into a form or a file
 - PUT : upload a representation of the specified resource
 - DELETE : delete the specified resource
- relative part of URI: (contains, e.g., path/object) and identifies a resource relative to server
- HTTP version:
 - HTTP/1.1 for HTTP 1.1 ^[2]



- HTTP method:
 - GET : download the data that belongs to a specified resource (can include parameters), e.g., a web page
 - HEAD : returns only the headers of a GET response (no request body)
 - POST : data is sent in the request body that should be handed to the requested resource, maybe the data filled into a form or a file
 - PUT : upload a representation of the specified resource
 - DELETE : delete the specified resource
- relative part of URI: identifies a resource relative to server
- HTTP version:
 - HTTP/1.1 for HTTP 1.1 ^[2]
 - HTTP/1.0 for HTTP 1.0 ^[1]

header field name : field value \r \n

- Different header fields, each has its own format

header field name : field value \r \n

- Different header fields, each has its own format
- Examples

header field name : field value \r \n

- Different header fields, each has its own format
- Examples
 - Accept: text/html

header field name : field value \r \n

- Different header fields, each has its own format
- Examples
 - Accept: text/html
 - Accept: image/jpg

header field name : field value \r \n

- Different header fields, each has its own format
- Examples
 - Accept: text/html
 - Accept: image/jpg
 - Accept-Language: zh-CN

header field name : field value \r \n

- Different header fields, each has its own format
- Examples
 - Accept: text/html
 - Accept: image/jpg
 - Accept-Language: zh-CN
 - Accept-Charset: utf-8, gb2312

header field name : field value \r \n

- Different header fields, each has its own format
- Examples
 - Accept: text/html
 - Accept: image/jpg
 - Accept-Language: zh-CN
 - Accept-Charset: utf-8, gb2312
 - If-Modified-Since: Wed, 27 Mar 2013 12:01:32 GMT

- Let's look how requests generated by web browsers look like

Listing: HTTPServerPrintingRequests.java

```
import java.io.BufferedReader;           import java.io.BufferedWriter;           import java.io.InputStreamReader;
import java.io.OutputStreamWriter;       import java.net.ServerSocket;           import java.net.Socket;

public class HTTPServerPrintingRequests {
    public static final void main(final String[] args) {
        String s;           StringBuilder sb;

        try (ServerSocket server = new ServerSocket(9995)) { //create server socket
            try (Socket client = server.accept()) { //accept incoming client

                sb = new StringBuilder(); //allocate buffer

                try (InputStreamReader ir = new InputStreamReader(client.getInputStream()); //request=character stream
                    BufferedReader br = new BufferedReader(ir)) { //read request line-by-line
                    while ((s = br.readLine()) != null) { //as long as lines can be read...
                        sb.append(s); //append them to the buffer
                        sb.append("<br/>"); //add HTML line breaks
                        if(s.length()<=0) { break; } // the final newline of the header
                    }
                    client.shutdownInput(); //no more input is requests

                    try (OutputStreamWriter pw = new OutputStreamWriter(client.getOutputStream())) {
                        pw.write("HTTP/1.1 200 OK\r\n\r\n<html><body><pre>"); //now write the answer: HTTP OK + HTML document
                        pw.write(sb.toString()); //buffered content
                        pw.write("</pre></body></html>"); //close the HTML document
                    }
                }
            }
        } catch (Throwable t) {
            t.printStackTrace();
        }
    }
}
```


We run the `HTTPServerPrintingRequests` locally and access `localhost:9995` with Firefox

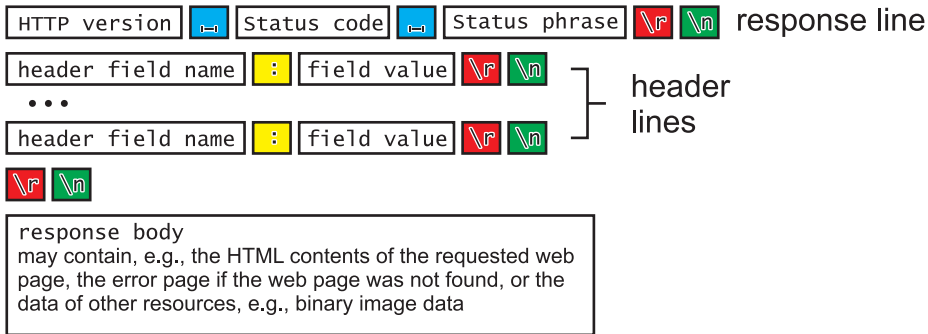
We run the HTTPServerPrintingRequests locally and access localhost:9995 with Firefox:

```
GET / HTTP/1.1
Host: localhost:9995
User-Agent: Mozilla/5.0 (Windows NT 6.1; WOW64; rv:19.0) Gecko/20100101 Firefox/19.0
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8
Accept-Language: en-us,en;q=0.8,de-de;q=0.5,de;q=0.3
Accept-Encoding: gzip, deflate
DNT: 1
Connection: keep-alive
Cache-Control: max-age=0
```

We run the `HTTPServerPrintingRequests` locally and access `http://localhost:9995` with Internet Explorer

We run the `HTTPServerPrintingRequests` locally and access `http://localhost:9995` with Internet Explorer:

```
GET / HTTP/1.1
Accept: application/x-ms-application, image/jpeg, application/xaml+xml, image/gif,...
Accept-Language: de-DE
User-Agent: Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 6.1; WOW64; Trident/4.0;;...
Accept-Encoding: gzip, deflate
Host: localhost:9995
Connection: Keep-Alive
```



HTTP version  Status code  Status phrase   response line

- HTTP version

HTTP version  Status code  Status phrase   response line

- HTTP version:
 - HTTP/1.1 for HTTP 1.1 ^[2]

HTTP version  Status code  Status phrase   response line

- HTTP version:
 - HTTP/1.1 for HTTP 1.1 ^[2]
 - HTTP/1.0 for HTTP 1.0 ^[1]

HTTP version  Status code  Status phrase   response line

- HTTP version:
 - HTTP/1.1 for HTTP 1.1 ^[2]
 - HTTP/1.0 for HTTP 1.0 ^[1]
- Status code

HTTP version  Status code  Status phrase   response line

- HTTP version:
 - HTTP/1.1 for HTTP 1.1 ^[2]
 - HTTP/1.0 for HTTP 1.0 ^[1]
- Status code: three digits, with the first digit representing

HTTP version  Status code  Status phrase   response line

- HTTP version:
 - HTTP/1.1 for HTTP 1.1 ^[2]
 - HTTP/1.0 for HTTP 1.0 ^[1]
- Status code: three digits, with the first digit representing

HTTP version  Status code  Status phrase   response line

- HTTP version:
 - HTTP/1.1 for HTTP 1.1 ^[2]
 - HTTP/1.0 for HTTP 1.0 ^[1]
- Status code: three digits, with the first digit representing ¹ Information

HTTP version  Status code  Status phrase   response line

- HTTP version:
 - HTTP/1.1 for HTTP 1.1 ^[2]
 - HTTP/1.0 for HTTP 1.0 ^[1]
- Status code: three digits, with the first digit representing
 - 1 Information
 - 2 Success

HTTP version  Status code  Status phrase   response line

- HTTP version:
 - HTTP/1.1 for HTTP 1.1 ^[2]
 - HTTP/1.0 for HTTP 1.0 ^[1]
- Status code: three digits, with the first digit representing
 - 1 Information
 - 2 Success
 - 3 Redirection

HTTP version  Status code  Status phrase   response line

- HTTP version:
 - HTTP/1.1 for HTTP 1.1 ^[2]
 - HTTP/1.0 for HTTP 1.0 ^[1]
- Status code: three digits, with the first digit representing
 - 1 Information
 - 2 Success
 - 3 Redirection
 - 4 Error on client side (e.g., wrong URI)

HTTP version  Status code  Status phrase   response line

- HTTP version:
 - HTTP/1.1 for HTTP 1.1 ^[2]
 - HTTP/1.0 for HTTP 1.0 ^[1]
- Status code: three digits, with the first digit representing
 - 1 Information
 - 2 Success
 - 3 Redirection
 - 4 Error on client side (e.g., wrong URI)
 - 5 Error on server side

HTTP version  Status code  Status phrase   response line

- HTTP version:
 - HTTP/1.1 for HTTP 1.1 ^[2]
 - HTTP/1.0 for HTTP 1.0 ^[1]
- Status code: three digits, with the first digit representing
 - 1 Information
 - 2 Success
 - 3 Redirection
 - 4 Error on client side (e.g., wrong URI)
 - 5 Error on server side
- Status phrase

HTTP version  Status code  Status phrase   response line

- HTTP version:
 - HTTP/1.1 for HTTP 1.1 ^[2]
 - HTTP/1.0 for HTTP 1.0 ^[1]
- Status code: three digits, with the first digit representing
 - 1 Information
 - 2 Success
 - 3 Redirection
 - 4 Error on client side (e.g., wrong URI)
 - 5 Error on server side
- Status phrase: short textual representation of status code, e.g., OK

- Some examples for status codes

- Some examples for status codes
 - 200 OK

- Some examples for status codes
 - 200 OK
 - 301 Moved Permanently

- Some examples for status codes
 - 200 OK
 - 301 Moved Permanently
 - 400 Bad Request

- Some examples for status codes
 - 200 OK
 - 301 Moved Permanently
 - 400 Bad Request
 - 404 Not Found

- Some examples for status codes
 - 200 OK
 - 301 Moved Permanently
 - 400 Bad Request
 - 404 Not Found
 - 505 HTTP Version Not Supported

header field name : field value \r \n

- Different header fields, each has its own format

header field name : field value \r \n

- Different header fields, each has its own format
- Examples

header field name : field value \r \n

- Different header fields, each has its own format
- Examples
 - Content-Type: text/html

header field name : field value \r \n

- Different header fields, each has its own format
- Examples
 - Content-Type: text/html
 - Content-Length: 16384

header field name : field value \r \n

- Different header fields, each has its own format
- Examples
 - Content-Type: text/html
 - Content-Length: 16384
 - Language: zh-CN;

header field name : field value \r \n

- Different header fields, each has its own format
- Examples
 - Content-Type: text/html
 - Content-Length: 16384
 - Language: zh-CN;
 - Last-modified: 28 Mar 2013

- Let's look how responses generated by web servers look like

Listing: MinHTTPClientJava17.java Min HTTP Client + Try-With-Resource

```
import java.io.BufferedReader;          import java.io.InputStreamReader;
import java.io.OutputStreamWriter;      import java.net.Socket;

public class MinHTTPClientJava17 {//this is a minimum web client; see lesson 07 coming later
    public static final void main(final String[] args) {
        String dest, request, response;

        dest    = "www.baidu.com";      // a random example for a Chinese host
        request = "GET_/index.htmlHTTP/1.1\nHost:_" + dest + "\n\n";

        try(Socket sock = new Socket(dest, 80)) { // web servers are usually listening at port 80
            try(OutputStreamWriter w = new OutputStreamWriter(sock.getOutputStream())) {
                w.write(request);          // write the HTTP request ^-^
                w.flush();                // make sure that all data has been sent
                sock.shutdownOutput();    // closing down the channel for sending data to the server

                try (InputStreamReader is = new InputStreamReader(sock.getInputStream());
                    BufferedReader r = new BufferedReader(is)) { // Baidu uses UTF-8 encoding
                    while ((response = r.readLine()) != null) { // read strings line-by-line until connection closed by server
                        System.out.println(response);          // print to output
                    }
                }
            }
        } catch (Throwable t) {
            t.printStackTrace();
        }
    }
}
```


To `www.baidu.com`, we send

To `www.baidu.com`, we send:

```
GET /index.html HTTP/1.1  
Host: www.baidu.com
```

To `www.baidu.com`, we send:

```
GET /index.html HTTP/1.1
Host: www.baidu.com
```

and get the response:

```
HTTP/1.1 200 OK
Date: Wed, 27 Mar 2013 23:44:43 GMT
Server: BWS/1.0
Content-Length: 10319
Content-Type: text/html;charset=utf-8
Cache-Control: private
Expires: Wed, 27 Mar 2013 23:44:43 GMT
Set-Cookie: H_PS_PSSID=2097_1430_2132_1945_1788; path=/; domain=.baidu.com
Set-Cookie: BAIDUID=1BBB7C987D5159BE0741B675A88B3E0C:FG=1; expires=Wed, 27-Mar-43...
P3P: CP=" OTI DSP COR IVA OUR IND COM "
Connection: Keep-Alive

<!DOCTYPE html><!--STATUS OK--> <html><head> <meta http-equiv="content-type"
content="text/html;charset=utf-8"> <title>...
```

- HTTP is a general, text-based protocol to request resources.
- Web pages are served by web servers which implement the HTTP protocol.
- Such servers can be implemented with the simple stuff we have learned about sockets.
- Actually, several Java-based web servers use exactly the technologies we have learned so far.

谢谢

Thank you

Thomas Weise [汤卫思]
tweise@hfu.edu.cn
<http://www.it-weise.de>

Hefei University, South Campus 2
Institute of Applied Optimization
Shushan District, Hefei, Anhui,
China



Caspar David Friedrich, "Der Wanderer über dem Nebelmeer", 1818
http://en.wikipedia.org/wiki/Wanderer_above_the_Sea_of_Fog



1. Timothy John Berners-Lee, R. Fielding, and H. Frystyk. *Hypertext Transfer Protocol – HTTP/1.0*, volume 1945 of *Request for Comments (RFC)*. Network Working Group, May 1996. URL <http://tools.ietf.org/html/rfc1945>.
2. R. Fielding, J. Gettys, Jeffrey Mogul, H. Frystyk, L. Masinter, P. Leach, and Timothy John Berners-Lee. *Hypertext Transfer Protocol – HTTP/1.1*, volume 2616 of *Request for Comments (RFC)*. Network Working Group, June 1999. URL <http://tools.ietf.org/html/rfc2616>.
3. David Gourley and Brian Totty. *HTTP: The Definitive Guide*. Definitive Guide. Sebastopol, CA, USA: O'Reilly Media, Inc., 2002. ISBN 1565925092 and 9781565925090. URL http://books.google.de/books?id=qEo019bcV_cC.