

Exercise Session 3: Peeking under the Web

COM-208: Computer Networks

The main goal of this exercise session is to get a sense of how web browsers and web servers communicate: we will play around with **caching** and **cookies**.

Caching at the web browser

Web browsers **cache** resources, so that they don't need to download them again if the user requests them again. You will now experience the difference this browser behavior can make.

Use a Firefox web browser, if you can. It comes with a nice tool, the web-developer network console ($\equiv \rightarrow$ **More tools** \rightarrow **Web Developer Tools** \rightarrow **Network**), which visualizes each HTTP request that the browser makes, as well as the corresponding HTTP response that the browser receives. If you click on an HTTP request from the list on the left, you will see all the relevant information in the panel on the right.

Get ready to capture web traffic:

- Open your web browser and clear the cache. To do so in Firefox: $\equiv \rightarrow$ **Settings** \rightarrow **Privacy & Security** \rightarrow **Cookies and Site Data** \rightarrow **Clear Data...**
- In Firefox, open the web-developer network console.
- Open Wireshark and start a new traffic capture.

The resources in this exercise are downloaded from web servers that use HTTPS instead of HTTP. HTTPS is a secure version of HTTP: HTTP messages encrypted within SSL (the Secure Sockets Layer that we mentioned in class) packets. This makes using Wireshark a bit harder, but we will guide you:

- To find out the web servers your web browser sent requests to, apply the filter: "**ssl and ssl.handshake.type==1**". This will display all the **Client Hello** messages. A **Client Hello** is the first packet that your computer sent to the web server to initiate their communication.

- To see the rest of packets exchanged with the web server: select the `Client Hello` message of that server, go to `Analyze` → `Follow` → `TCP Stream`. `Ignore/close` the window that pops up. Now you should see only the packets that belong to the same TCP connection as the packet you chose.

Answer the following questions, using the web-developer network console, or Wireshark, or (ideally) both.

- Visit [Welcome to Rio](#). Where is this resource downloaded from?
- How long did it take to download it?
- Restart your web browser. Visit [Welcome to Rio](#) again. How long did it take to load it this time? What explains the difference?
- Open a second tab in your web browser and visit [Welcome to Rio II](#). Where is this resource downloaded from? How long did it take to download it?
- When you visited ‘Welcome to Rio II’, your web browser had already cached ‘Welcome to Rio’, which is essentially the same image. Do you think your browser served ‘Welcome to Rio II’ from the cache, or it downloaded it from its origin web server? Why do you think your browser behaved this way?

Caching at a proxy web server

It is not only web browsers that cache resources; **proxy web servers** are web servers that act as **intermediaries**: they cache resources that are originally stored in other web servers (called **origin web servers**) and serve them to nearby web clients.

Before you start, clear your browser cache and find the proxy settings of your web browser. In Firefox: $\equiv \rightarrow$ Settings \rightarrow General \rightarrow Network Settings \rightarrow Settings. Setup a proxy web server using the following settings:

- Check “Manual proxy configuration”.
- Set HTTP Proxy: 172.105.184.208 and Port: 8001 (you could use any proxy web server, that does HTTPS caching, from <https://free-proxy-list.net/>).
- Check “Also use this proxy for HTTPS”.

Now visit the same two resources that you visited before.

- Where were the resources downloaded from?
- How long did it take to download each resource this time? Why did the download time change?
- What benefit would a network (e.g., EPFL network) gain by installing a proxy web server?
- What will happen to your web browser if the proxy web server that you specified fails? Will your browser be able to load any new web pages? What about pages that are cached?

IMPORTANT: Restore your original proxy settings: $\equiv \rightarrow$ Settings \rightarrow General \rightarrow Network Settings \rightarrow Settings \rightarrow Use system proxy settings.

Cookies

Cookies enable a web server to **link subsequent HTTP requests** to the same web browser: if you send 10 HTTP GET requests, for 10 different resources, to the same web server, the web server can use cookies to figure out that these 10 requests came from the same web browser, even if you did not explicitly provide any identification information (e.g., you did not login).

Before you start, figure out how to control cookie settings in your browser. In Firefox:

- To view or delete the cookies that have been stored on your computer: $\equiv \rightarrow$ **Settings** \rightarrow **Privacy & Security** \rightarrow **Cookies and Site Data** \rightarrow **Manage Data** or **Clear Data...**
- You can also view the cookies that your computer sends along with an HTTP request, or receives along with the corresponding response, through the web developer network console: select an HTTP request from the list of requests on the left, then select the **Cookies** menu from the panel on the right.
- To view all the cookies stored due to visiting the web page: Go to **Storage** from the top panel, and then select **Cookies**.

Let us see cookies in action:

- Allow your browser to exchange cookies. Delete existing cookies. Visit [MeteoSuisse](#). Did the MeteoSuisse web server send you any cookies?
- By default, MeteoSuisse shows you the weather for Geneva. Choose another location for which you want to see the weather. Restart your web browser and re-visit [MeteoSuisse](#). Do you get the weather for Geneva as before? Explain your browser's behavior.
- Delete existing cookies. Restart your web browser and re-visit [MeteoSuisse](#). Do you get the weather for Geneva or for your chosen location? Explain your browser's behavior.