

Wesleyan University, Fall 2022, COMP 411

Homework 4: DNS and Web Proxies

Due by 5pm on October 7, 2022

1. WRITTEN PROBLEMS (5 POINTS)

PROBLEM 1. *This question will use the command-line tool `dig` to explore the DNS hierarchy of servers. To learn more about how `dig` works, type the command `man dig` in a terminal. Starting with a root DNS servers (e.g., choose one from <https://www.iana.org/domains/root/servers>), you will initiate a sequence of non-recursive DNS look-ups to resolve a hostname using `dig`. To help you get started, an example first command might be the following, using `a.root-servers.net` as the name server. Note the use of the `@` sign in the command.*

`dig +norecurse @a.root-servers.net any www.wesleyan.edu` *This is for http website*
`dig +norecurse @a.root-servers.net any wesleyan.edu` *This is for https website*
From this command you'll get a list of next name servers. Choose one of the name servers and execute another `dig` command. Repeat this process until the address is resolved.

(a): *Give the list of all name servers you use to resolve `www.wesleyan.edu`. List the name servers in the order that you accessed them. Describe anything interesting that you notice.*

(b): *Give the list of all name servers you use to resolve the host `www.csail.mit.edu`. List the name servers in the order that you accessed them.*

Solution:

(a): The following name servers were used to resolve `www.wesleyan.edu`.

```
a.root-servers.net
f.edu-servers.net
ns3.wesleyan.edu
```

First command: `dig +norecurse @a.root-servers.net any www.wesleyan.edu`
First response contained the following name servers:

```
;; AUTHORITY SECTION:
edu. 172800 IN NS b.edu-servers.net.
edu. 172800 IN NS f.edu-servers.net.
edu. 172800 IN NS i.edu-servers.net.
edu. 172800 IN NS a.edu-servers.net.
edu. 172800 IN NS g.edu-servers.net.
edu. 172800 IN NS j.edu-servers.net.
edu. 172800 IN NS k.edu-servers.net.
edu. 172800 IN NS m.edu-servers.net.
```

```
edu. 172800 IN NS l.edu-servers.net.
edu. 172800 IN NS h.edu-servers.net.
edu. 172800 IN NS c.edu-servers.net.
edu. 172800 IN NS e.edu-servers.net.
edu. 172800 IN NS d.edu-servers.net.
```

Second command: `dig +norecurse @b.edu-servers.net any www.wesleyan.edu`
 Second response contained the following name servers:

```
;; AUTHORITY SECTION:
wesleyan.edu. 172800 IN NS aws-ns0.wesleyan.edu.
wesleyan.edu. 172800 IN NS aws-ns1.wesleyan.edu.
wesleyan.edu. 172800 IN NS aws-ns2.wesleyan.edu.
wesleyan.edu. 172800 IN NS aws-ns3.wesleyan.edu.
wesleyan.edu. 172800 IN NS ns0.wesleyan.edu.
wesleyan.edu. 172800 IN NS ns1.wesleyan.edu.
```

Third command: `dig +norecurse @aws-ns3.wesleyan.edu any www.wesleyan.edu`
 Third response contained the following answer:

```
;; ANSWER SECTION:
www.wesleyan.edu. 600 IN CNAME www-load-balancer-1666265967.us-east-1.elb.amazonaws.com.
```

This tells us that to access the Wesleyan website, we go through a load balancer managed by Amazon AWS. We must therefore do a dig on this name.

Fourth command: `dig +norecurse @a.root-servers.net any www-load-balancer-1666265967.us-`
 Fourth response contained the following answer:

```
;; AUTHORITY SECTION:
com. 172800 IN NS e.gtld-servers.net.
com. 172800 IN NS b.gtld-servers.net.
com. 172800 IN NS j.gtld-servers.net.
com. 172800 IN NS m.gtld-servers.net.
com. 172800 IN NS i.gtld-servers.net.
com. 172800 IN NS f.gtld-servers.net.
com. 172800 IN NS a.gtld-servers.net.
com. 172800 IN NS g.gtld-servers.net.
com. 172800 IN NS h.gtld-servers.net.
com. 172800 IN NS l.gtld-servers.net.
com. 172800 IN NS k.gtld-servers.net.
com. 172800 IN NS c.gtld-servers.net.
com. 172800 IN NS d.gtld-servers.net.
```

Fifth command: `dig +norecurse @e.gtld-servers.net any www-load-balancer-1666265967.us-e`
 Fifth response contained the following answer:

```
;; AUTHORITY SECTION:
amazonaws.com. 172800 IN NS ns-27.awsdns-03.com.
amazonaws.com. 172800 IN NS ns-967.awsdns-56.net.
amazonaws.com. 172800 IN NS ns-1670.awsdns-16.co.uk.
amazonaws.com. 172800 IN NS ns-1321.awsdns-37.org.
```

Sixth command: `dig +norecurse @ns-1119.awsdns-11.org any www-load-balancer-1666265967.us-east-1.elb.amazonaws.com`
 Sixth response contained the following answer:

```
;; ANSWER SECTION:
www-load-balancer-1666265967.us-east-1.elb.amazonaws.com. 60 IN A 3.217.226.105
www-load-balancer-1666265967.us-east-1.elb.amazonaws.com. 60 IN A 3.213.45.230
```

This is the address for the http version of the Wesleyan webpage (you will get a warning but okay to click through).

To go to the https version of the Wesleyan webpage, we do a dig on wesleyan.edu as follows.

First command: `dig +norecurse @a.root-servers.net any wesleyan.edu`

First response contained the following name servers:

```
;; AUTHORITY SECTION:
edu. 172800 IN NS b.edu-servers.net.
edu. 172800 IN NS f.edu-servers.net.
edu. 172800 IN NS i.edu-servers.net.
edu. 172800 IN NS a.edu-servers.net.
edu. 172800 IN NS g.edu-servers.net.
edu. 172800 IN NS j.edu-servers.net.
edu. 172800 IN NS k.edu-servers.net.
edu. 172800 IN NS m.edu-servers.net.
edu. 172800 IN NS l.edu-servers.net.
edu. 172800 IN NS h.edu-servers.net.
edu. 172800 IN NS c.edu-servers.net.
edu. 172800 IN NS e.edu-servers.net.
edu. 172800 IN NS d.edu-servers.net.
```

Second command: `dig +norecurse @b.edu-servers.net any wesleyan.edu`

Second response contained the following name servers:

```
;; AUTHORITY SECTION:
wesleyan.edu. 172800 IN NS aws-ns0.wesleyan.edu.
wesleyan.edu. 172800 IN NS aws-ns1.wesleyan.edu.
wesleyan.edu. 172800 IN NS aws-ns2.wesleyan.edu.
wesleyan.edu. 172800 IN NS aws-ns3.wesleyan.edu.
wesleyan.edu. 172800 IN NS ns0.wesleyan.edu.
wesleyan.edu. 172800 IN NS ns1.wesleyan.edu.
```

Third command: `dig +norecurse @aws-ns0.wesleyan.edu any wesleyan.edu`
 Third response contained the following answer:

```
;; ANSWER SECTION:
wesleyan.edu. 14400 IN A 52.204.12.162
wesleyan.edu. 14400 IN NS aws-ns0.wesleyan.edu.
wesleyan.edu. 14400 IN NS aws-ns1.wesleyan.edu.
wesleyan.edu. 14400 IN NS aws-ns2.wesleyan.edu.
wesleyan.edu. 14400 IN NS aws-ns3.wesleyan.edu.
wesleyan.edu. 14400 IN NS ns0.wesleyan.edu.
wesleyan.edu. 14400 IN NS ns1.wesleyan.edu.
wesleyan.edu. 14400 IN SOA ns.wesleyan.edu. system.wesleyan.edu. 108406611 3600 300 2419
wesleyan.edu. 600 IN MX 10 usb-smtp-inbound-1.mimecast.com.
wesleyan.edu. 600 IN MX 10 usb-smtp-inbound-2.mimecast.com.
wesleyan.edu. 14400 IN TXT "3k49rknfl8au6ckcrij72cesc2"
wesleyan.edu. 14400 IN TXT "MS=B39DB878794B6ECFC79F2E381A25CFBDA8D6BA0F"
wesleyan.edu. 14400 IN TXT "ZOOM_verify_epShMhxwRNSeNnbdFsZHbQ"
wesleyan.edu. 14400 IN TXT "adobe-idp-site-verification=a116017a-f3c5-48e0-a6b9-1c16d913
wesleyan.edu. 14400 IN TXT "apple-domain-verification=sEQuS0b9KOE7Zdzq"
wesleyan.edu. 14400 IN TXT "citrix-verification-code=ff713e90-387d-4840-a02e-b2fb8fa4d4e
wesleyan.edu. 14400 IN TXT "google-site-verification=Ttbs_u7D1M1m8yzfxk4PSehjhr9kQ9LivVf
wesleyan.edu. 14400 IN TXT "v=spf1 redirect=_spf.wesleyan.edu"
wesleyan.edu. 14400 IN TXT "y11w4D0fonmrCSC9qQHCsFAXwkuiJfQmLcaaipgapE5044mydlAKQ/1yeT31"
```

Third response contained the IP address for `wesleyan.edu`.

(b): The following name servers were used to resolve `www.csail.mit.edu`.

```
dig +norecurse @a.root-servers.net any csail.mit.edu
dig +norecurse @f.edu-servers.net any csail.mit.edu
dig +norecurse @usw2.akam.net any csail.mit.edu
dig +norecurse @auth-ns3.csail.mit.edu any csail.mit.edu
csail.mit.edu. 3600 IN A 128.30.2.109
```

2. CODING (15 POINTS)

PROBLEM 2. In Homework 3, you created a simple web client and web proxy (Part 1). In this homework you will extend your programs to include a cache (Part 2), use conditional GET requests (Part 3), and work with a real web browser (Part 4). For reference the web client and web proxy interactions are again as in the setup shown in Figure 1.

Part 2: Proxy cache. Part 2 is an intermediate step on the way to making the web proxy more realistic. The purpose of Part 2 is simply to store the responses from the various web servers to which clients connect.

What to cache? An HTTP response message has the format below. You may additionally see the header line, `Last-Modified:<date>` in the HTTP response. In a real web proxy, you would

I'd like www.nytimes.com...

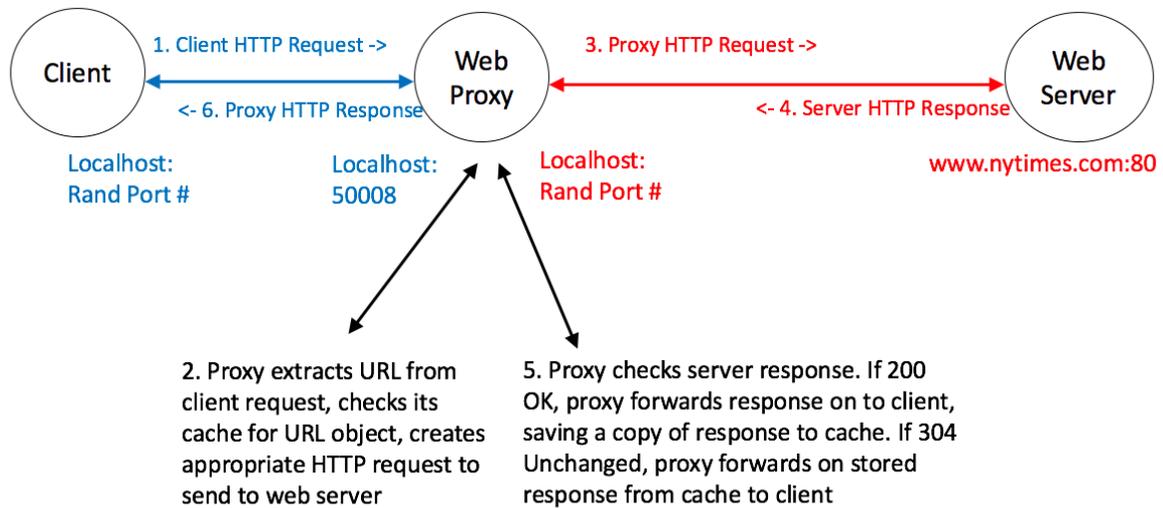


FIGURE 1. Architecture for Problem 3 programming.

store only the entity-body along with the URL and date. However, for simplicity, and so that you don't have to parse the HTTP responses, you will save the entire HTTP response in your cache. You will want to keep track of the HTTP response itself, the URL it is associated with, and the Last-Modified date returned in the response indicating when it was last updated (or simply the Date value if there is no Last-Modified field). This information will be used in Part 3 of the project.

```
HTTP/1.1 200 OK\r\n
Date: Tue, 20 Jan 2018 18:27:37 GMT \r\n
Connection: Close \r\n
Content-Type: text/plain \r\n
Content-Length: 5430 \r\n
Last-Modified: Sun, 18 Jan 2018 20:43:27 GMT\r\n\r\n
[ Entity-body of length 5430 bytes ]
```

Cache data structure. You are free to choose whatever approach you'd like to organize your cache, which might range from a simple dictionary data structure, to storing webpages in files. I suggest starting with something very simple like a dictionary, so that you can move on to Part 3.

Part 3: Conditional GET requests. Rather than have the web proxy automatically forward on the client's GET request to the web server, the web proxy will now forward on a modified version of the GET request. How the request is modified depends on whether the webpage object for the desired URL is stored in the proxy cache. If the object for the desired URL is not stored in the proxy cache, the web proxy forwards on the client's GET request unmodified. If the object for the desired URL is stored in the proxy cache, the web proxy turns the client's GET request into

a conditional *GET* request, using the date information stored in the cache. A conditional *HTTP GET* request has the following format:

```
GET /mathcs/index.html HTTP/1.1\r\n
Host: www.wesleyan.edu\r\n
If-Modified-Since: Sun, 18 Jan 2018 20:43:27 GMT\r\n\r\n
```

In the response to the conditional *GET* request, one of two things may occur:

- (1) **The URL object has changed, or the If-Modified-Since is ignored.** In this case, the web server's response will be as in Part 1, and the web proxy forwards the web server's *HTTP* response directly to the client.
- (2) **The URL object has not changed.** The web server will return a 304 Not Modified status code in its response rather than a 200 OK status code. In this case what the web proxy has stored in its cache is what is returned to the client. The web proxy creates an *HTTP* response containing the information in its cache, and forwards this response to the client. If the web proxy is storing entire responses in its cache, rather than just the entity-body, as you will be doing, then the web proxy simply returns what is stored in its cache, which should be a validly formatted *HTTP* Response.

Several websites you might wish to use to test conditional requests include `www.mit.edu` and `www.cyberciti.biz`. I recommend also testing with a URL which includes a pathname and object such as `www.wesleyan.edu/mathcs/index.html`.

Part 4: Testing with a web browser. In Part 4, you will test your web proxy using a web browser as a client instead of your web client. To do this, choose your favourite browser and change the proxy settings to use the web proxy listening on `localhost:50007`: i.e., to use your web proxy.

Pathname parsing. When a web browser uses a proxy, it includes the full URL including the hostname as the path in the *GET* line of the request. This pathname, however, will not be parseable by web servers. Therefore, the web proxy must now check the pathname it is given, identify whether it is a complete URL, and if so create a new *HTTP* request that includes only the path and object portion of the URL in its *GET* line.

You may find the following helpful when looking at *HTTP* requests and responses.

- **RFC 2616:** Hypertext Transfer Protocol – HTTP/1.1. <https://tools.ietf.org/html/rfc2616>
- **RFC 7232:** Hypertext Transfer Protocol (HTTP/1.1): Conditional Requests. <https://tools.ietf.org/html/rfc7232>

For this problem you should submit the following files. I recognize that as this is a two-part programming assignment, that things such as the design may change in the second part. However, part of my goal is to get you think about how you design, structure, and test your code.

- *All of your python files.* Your python code should include a header at the top of each file describing what the file does and your python code itself should be well-commented, check for exceptions and shut down nicely (e.g., clean up sockets and other state when the web client or web proxy terminates for some reason).
- *Readme file* indicating how to run your code and listing all of your python files and the purpose of each file.
- *A write-up* describing and motivating your program design, a list of web sites for which your program works, a list of example websites for which your program does not work (if any) along with an explanation why. For at least one test case for which your program works, include example screen shots of the output. I do not expect that your code will work with all web servers, since, for instance, some web servers do not use utf-8 encoding, and you are not fully parsing responses. However, in such cases you should catch whatever error is shown and continue cleanly or shutdown. Also include any test cases that cause the web client or web proxy to crash.

3. SUBMISSION

Upload your written work as **hw4.pdf**, all files (***.py, README, and your design**) to the Google Drive directory I have created for you named `comp411-f22-USERNAME/hw4/`. You should replace `USERNAME` with your Wesleyan username.

Do not forget that your written work must be submitted as a PDF! Make sure that at the top of each file you have put your name! Do not, however, change the names of the files.