15-441/641: Computer Networks
Project 1: A Web Server Called Liso

TAs: Krithika Vijayan <kvijayan@andrew.cmu.edu>
Alex Bainbridge <abainbri@andrew.cmu.edu>

Assigned: January 18, 2019
Checkpoint 1 due: January 25, 2019
Checkpoint 2 due: February 8, 2019
Final version due: February 15, 2019

1 Introduction

The purpose of this project is to give you experience in designing and developing concurrent network applications. You will use the Berkeley Sockets API to write a web server using a subset of the HyperText Transport Protocol (HTTP) 1.1 Request for Comment—RFC 2616 [2]. Your web server will also implement HyperText Transport Protocol Secure (HTTPS) via Transport Layer Security (TLS) as described in RFC 2818 [3]. The final part of the project will implement the Common Gateway Interface (CGI) as described in RFC 3875 [4]. This set of features forms the core of the Liso web server's capabilities. Through this project, you will have the opportunity to synthesize the provided specifications into a carefully designed implementation.

Why are we doing this? First of all, we want you to get started with a big project in C – and this project should be familiar to you after implementing the proxy lab for 15-213/513. Second of all, most of us think of web services when we think of using the Internet – but there are a lot of moving parts (routing, DNS, reliable transport) that web servers rely on from networking to do their job correctly. This project will get you thinking about all of the tasks that networking does for you automatically just to make your web server do its job. The next few weeks of lectures should remove the mystery/magic of what the network does 'underneath' your web server.

In building this project, you will become familiar with HTTP, HTTPS, and CGI. Your Liso web server will be fully functional and capable of running interactive applications via its CGI interface. At the end of the a project the final test will be running a simple python script to generate your own ASCII Art!

Be prepared: this is a single person project and it has a lot of depth—your skills will be exercised (perhaps to their limits). So start early and feel more than welcome to ask questions. The lead TAs on this project are Alex and Krithika and you can ask them for help during their office hours, but you can of course use the office hours of the other TAs as well. Note that the TAs are not allowed to debug your code for you. Specifically, they will only look at your code for a limited time (up to 10 minutes; leaving them time to help other students) and they will not modify your code (they cannot touch the keyboard). Their role

is to help you learn how to debug, and to answer any questions you have about the project.

2 The Liso Server

Your server will implement HEAD, GET (both needed for HTTP 1.1 general purpose server compliance), and POST. This should comply with the specification in the RFC [2]. After this we will move on into implementing TLS, and finally CGI.

2.1 Supporting HTTP 1.1

- **GET** – requests a specified resource; it should not have any significance other than retrieval
- **HEAD** – asks for an identical response as GET, without the actual body—no bytes from the requested resource
- **POST** – submit data to be processed to an identified resource; the data is in the body of this request; side-effects expected

For all other commands, your server must return "501 Method Unimplemented." If you are unable to implement one of the above commands (perhaps you ran out of time), your server must return the error response "501 Method Unimplemented," rather than failing silently (or not so silently). While you develop, you may want to just return this error response always until features are implemented—no matter what you will have a valid HTTP 1.1 server!

Your server should be able to support multiple clients concurrently. The only limit to the number of concurrent clients should be the number of available file descriptors in the operating system (the min of ulimit -n and FD_SETSIZE—both typically 1024). We will not be testing beyond 1024 concurrent connections. While the server is waiting for a client to send the next command, it should be able to handle inputs from other clients. Also, your server should not hang if a client sends only a partial request. In general, concurrency can be achieved using either `select()` or multiple threads. However, in this project, you **must implement your server using `select()` to support concurrent connections**. Threads are **NOT** permitted at all for the project. See the resources section below for help on these topics. As a public server, your implementation should be robust to client errors. For example, your server should be able to handle malformed requests which do not have proper [CR][LF] line endings. The provided lex and yacc parser is not robust to these malformed requests and you are expected to add the necessary rules if you decide to go-ahead and use the provided parser. For example, it must not overflow any buffers when a malicious client sends a message that is "too long." These are some things we will test for.

You are implementing a *real*, standards-compliant web server. Therefore, comparing protocol exchanges to existing web servers is both valid and encouraged. Install Apache [5], install Wireshark [6], and sniff the protocol exchanges and compare them to your own—even use captured web browser requests to replay from files you save as input to your implementation of Liso. Come up with other create ways to test as well as there is a portion of the grade reserved for testing.

To aid you in programming, and testing, we have prepared a starter package available in the Autolab handout. It contains the starter code for an echo server and HTTP parser. This code needs to be modified to use select() as well as adding support for multiple clients at once. Also, the parsing rules need to be enhanced to support the RFC specification [2]. Additional information can be found in the checkpoint details Section 4.1.

2.2 HTTPS and CGI

In Checkpoint 3 you will extend your basic Liso server with support for HTTPS and CGI. You will use the OpenSSL library for HTTPS support. Specially, you will wrap communication calls with with SSL wrapping functions that will encrypt data sent over the channel and authenticate the server based on a certificate. The details of how this works will be covered later in the course when we talk about network security. In order to test HTTPS, you will need to get a certificate that your HTTPS-enabled server can use to authenticate itself to clients. The details are explained later in Section 4.3.

The Common Gateway Interface (CGI) provides a standard interface that allows a web server to call other servers. Web servers use this typically to generate dynamic content, i.e., the content that is generated is based on input provided by the client (using POST) or other client-specific information. More information on this part of the project is also provided in Section 4.3.

3 Guidelines for Implementation and Submission

This is a solo project: you *must* implement and submit your own code. Source materials for this project may be found on Autolab.

3.1 Coding and Version Control

Your server must be written in the C programming language. You are not allowed to use any custom socket classes or libraries, only the standard socket library and the provided library functions. You **may not** use the csapp wrapper library from 15-213/15-513 or libpthread for threading. We disallow csapp.c for two reasons: first, to ensure that you understand the raw standard BSD sockets API, and, second, because csapp.c's wrapper functions are not suitable for robust servers. Temporary system call failures (e.g., EINTR) in functions such as `select()` could cause the server to abort and utility functions like `rio_readlineb` are not designed for nonblocking code.

That said, we encourage the use of *anything* for testing. Use Wireshark [6], use telnet, use real web browsers, use Python to script tests—for testing, the sky is the limit.

All of your project files and submissions **must** be stored in a git repository.

You will submit your code as a tarball named `<andrewID>.tar`. Untarring this file should give us a directory named `15-441-project-1` which should contain the git repository as well as the code. You will submit this tarball using Autolab (`https://autolab.andrew.cmu.edu`). If you still can't login with your andrew ID by the end of January 22nd, let us know ASAP.

Checkpoints are designed to ensure that you keep tabs on your progress and are a great guideline to help you complete your project on time. Please note that Checkpoint 1 is

fairly straightforward and a small part of the total work. It helps you to familiarize yourself with the basics of socket programming (specifically how to use `select()` system call) and will require some amount of RFC interpretation for implementing a basic HTTP 1.1 parser. The rest of the checkpoints are progressively harder and will build on the previous checkpoints.

3.2 Compiling

You are responsible for making sure your code compiles and runs correctly on the Andrew x86 machines running Linux (i.e., `linux.andrew.cmu.edu` / `unix.andrew.cmu.edu`). We recommend using `gcc` to compile your program and `gdb` to debug it. You should use the `-Wall` and `-Werror` flags when compiling to generate full warnings and to help debug. Other tools available on the Andrew unix machines that are suggested are ElectricFence [11] (link with `-lefence`) and Valgrind [14]—use this with full leak checking to ensure you have no memory leaks. For this project, **you will also be responsible for turning in a GNU Make compatible Makefile**. See the GNU make manual[9] for details. When we run `make` we should end up with the Liso web server binary `lisod`.

3.3 Command Line Arguments

Liso will always have 8 arguments—functional or not(It may not be for the initial checkpoints):

usage: *./lisod <HTTP port> <HTTPS port> <log file> <lock file> <www folder> <CGI script path> <private key file> <certificate file>*

HTTP port – the port for the HTTP (or echo) server to listen on

HTTPS port – the port for the HTTPS server to listen on

log file – file to send log messages to (debug, info, error)

lock file – file to lock on when becoming a daemon process

www folder – folder containing a tree to serve as the root of a website

CGI script name (or folder) – for this project, this is a file that should be a script where you redirect all /cgi/* URIs. In the real world, this would likely be a directory of executable programs.

private key file – private key file path

certificate file – certificate file path

3.4 Running

The Liso server will be passed the ports to run on, what log file to use, what lock file to use when daemonizing, folders to serve static data from as well as CGI applications, and TLS private/public key pairs.

Not all of these options need to be functional at each stage of development. Only a HTTP port is needed for the first checkpoint when implementing an echo server using `select()`.

3.5 Framework Code

We will provide you with framework code that will, for example, help in forking a process for proper CGI handling and setting up the environment, parse commandline arguments (and sanity check them) and daemonize a process.

DISCLAIMER: We reserve the right to change the support code as the project progresses to fix bugs and to introduce new features that will help you debug your code. You are responsible for checking Piazza to stay up-to-date on these changes. We will assume that all students in the class will read and be aware of any information posted to Piazza.

3.6 HTTP Packet Parsing

Historical evidence suggests that most students spend considerable amount of time writing correct parsers. While parsing packets using C's string manipulation functions may well be an essential skill to have, it might get insanely tedious. We want you to spend time on other more important programming aspects such as socket programming, handling race conditions and memory leaks. For this reason, we require you to use Lex and Yacc for parsing packets. We will also provide you with a basic HTTP parser written in lex and yacc. More about parsing using Lex and Yacc will be covered in recitations. So, stay tuned!

4 Checkpoint Information

4.1 Checkpoint 1

4.1.1 Suggested Tasks

1. Unpack the starter code and create a git repo (`tar -zxvf Project1_starter.tar.gz; cd 15-441-project-1; git init`).

2. Create a `select()`-based echo server handling multiple clients at once, building on the starter code provided on autolab and the course website.

3. Specifically, your server must have the capability to parse HTTP 1.1 requests and classify them as "good" or "bad" based on the provided RFC [2]. For all "good" requests, you will simply echo back the original request. For all "bad" requests, you will return an HTTP response with 400 as the error code.

4. Test using our provided cp1_checker.py test script (read that script and understand it too.)

5. Finally, hand-in your submission by the deadline and include all needed files as outlined in §6.

To aid you in programming an echo server, and testing it, we have prepared a starter package for you available on Autolab. The server code needs to be modified to use select() as well as adding support for multiple clients. The lex and yacc starter code needs additional rules to do correct parsing.

4.2 Checkpoint 2

4.2.1 Suggested Tasks

1. Begin with your repository for Checkpoint 2, using the result of Checkpoint 1 as a starting point.
2. Create a simplified logging module for your project that writes out formatted logs. The format should begin with a header of the application name (Liso), the date, and the pid.
3. Make API functions for interacting with the logging module.
4. Log IP addreses, requests, and all errors such as in Common Log Format. [12]
5. Enhance your server to respond properly to any HTTP 1.1 request and implement persistent connections with HEAD, GET, and POST working as defined in RFC 2616. At this point as we don't have CGI we will not check responses to your POST requests but rather your ability to handle the request Body.
6. Ensure you read and write files from disk with full error handling (permissions, file does not exist, IO errors, etc.).
7. The server should also handle errors in a sane way. It should never completely crash (make it as robust as possible), send proper HTTP 1.1 error codes back to the browser and errors should be reported to a client as HTTP 1.1 error responses.
8. Submission is the same as Checkpoint 1. Tag and upload your repo in a tarball to the corresponding lab on Autolab.

4.2.2 Checkpoint 2 FAQ

1. Do we have to support pipelined requests?
   Yes
2. Do we have to support Chunking?
   No
3. For POST, if a Content-Length header is not provided in the request, is it fine to respond with a 411 as specified in the RFC?
   Yes, this is both possible and desired because it is a valid response. It may also simplify the design of the server.
4. Do we have to support "Conditional" GETs?
   No.
5. Can I use a hash table library written by another person?
   No, for this project implement your own if you really want it. You won't have to track every header, only the important ones for basic compliance.
6. Should we expect HTTP/1.1 requests to fit in one buffer?
   No, do not assume that they always fit inside one buffer. Be prepared to parse across buffer boundaries.

7. Can I assume that the request always has the Content-length field?
   Yes, assume requests to your server have Content-length if applicable. If it is missing, return a 411 response.

8. Since the server will serve static files, are we allowed to use http://svn.apache.org/repos/asf/httpd/httpd/trunk/docs/conf/mime.types
   Yes, You are allowed to use that or a simplified version of it.No need to support all well-known MIME-types, just the most common ones: text/html, text/css, image/png, image/jpeg, image/gif and maybe a few others up to your discretion.

9. For last-modified field, is there a C library that we can read file metadata such as this? Or is our web server supposed to handle that manually. (i.e. stamping each and every file in the www root with a last-modified stamp and manually updating that stamp every time a file is changed....).
   stat() is a system call to check for metadata on a file.

10. Are we allowed to reject requests with headers beyond a maximum size?
    Yes, you may reject any header line larger than 8192 bytes. Note, that this is different than a Content Length of greater than 8192. Additionally, you must find and parse the next request properly for pipelining purposes if you do reject the request.

11. Should we also handle request made up with /n instead of /r/n?
    No, some web servers do this and it is nice for telnet testing. However, Liso does not have this as a requirement.

4.3 Checkpoint 3

4.3.1 Suggested Tasks

1. Create support for the following CGI Variables. Not all variables may be used in the ASCII Art sample, but we will have additional checks to make sure that you do set these variables in case other CGI scripts are used. Using a python CGI script that echos the environment variables may be useful for your debugging.

   (a) CONTENT_LENGTH – taken directly from request
   (b) CONTENT_TYPE – taken directly from request
   (c) GATEWAY_INTERFACE – ”CGI/1.1”
   (d) PATH_INFO – $< path >$ component of URI
   (e) QUERY_STRING – parsed from URI as everything after “?”
   (f) REMOTE_ADDR – taken when accept() call is made
   (g) REQUEST_METHOD – taken directly from request
   (h) REQUEST_URI – taken directly from request
   (i) SCRIPT_NAME – hard-coded/configured application name (virtual path)
   (j) SERVER_PORT – as configured from command line (HTTP or HTTPS port depending)
   (k) SERVER_PROTOCOL – ”HTTP/1.1”
   (l) SERVER_SOFTWARE – ”Liso/1.0”

(m) HTTP_ACCEPT – taken directly from request
(n) HTTP_REFERER – taken directly from request
(o) HTTP_ACCEPT_ENCODING – taken directly from request
(p) HTTP_ACCEPT_LANGUAGE – taken directly from request
(q) HTTP_ACCEPT_CHARSET – taken directly from request
(r) HTTP_HOST – taken directly from request
(s) HTTP_COOKIE – taken directly from request
(t) HTTP_USER_AGENT – taken directly from request
(u) HTTP_CONNECTION – taken directly from request
(v) HTTP_HOST – taken directly from request

2. Additionally, your CGI module will need to do the following. Some good references: are the CGI RFC [4] and the URI RFC [1]. We have also provided a CGI runner in C and a runnable python script in the handout.

   (a) CGI URI's may accept GETs, POSTs, and HEADs; your job is not to decide this, just pass along information to the program being called
   (b) You must now parse inbound URI's to look for a special root folder, and in addition chop off arguments passed via URI according to specifications.
   (c) Any URI starting with "/cgi/" will be handled by a single command-line specified executable via a CGI interface coded by you
   (d) You need to pipe stdin, pipe stdout, fork(), setup environment variables per the CGI specification, and execve() the executable (it should be executable) Note: Watch the piped fd's in the parent process using your select() loop. Just add them to the appropriate select() sets and treat them like sockets, except you have to pipe them further to specific sockets.
   (e) Pass any message body (especially for POSTs) via stdin to the CGI executable
   (f) Receive any response over stdout until the process dies (monitor process status), or there is nothing more to read or broken pipe encountered
   (g) If the CGI program fails in any way, return a 500 response to the client, otherwise
   (h) Send all bytes from the stdout of the spawned process to the requesting client.
   (i) The CGI application will produce headers and message body as it sees fit, you do not need to modify or inspect these bytes at all; you are a proxy...things have come full circle!

3. Create a DNS hostname for yourself with a free account at No-IP [13] (or use a domain name you already have...)

4. In order to test your HTTPS implementation, we need to monitor your (encrypted!) HTTPS traffic. To allow this, add the 15-441 Carnegie Mellon University Root CA to your browser (import certificate, usually somewhere in preferences)

   (a) Now we can man-in-the-middle your HTTPS :-) Being course staff has perks!
   (b) But just trust us till this part is over...or make your own **CA**.

   (c) Really though, this is the part of the course where you need to Reflect on Trusting Trust.

5. Obtain your own private key and public certificate from the 15-441 CMU CA. You will need this when you run your HTTPS-enabled web server.

6. Implement SSL support - we have provided you a sample C server in the Autolab Handout.

   (a) Use the OpenSSL library. [7]

   (b) Create a second server listening socket in addition to the first one. Use the passed in SSL port from the commandline arguments.

   (c) Add this socket to the select() loop just like your normal HTTP server socket.

   (d) Whenever you accept connections, wrap them with the SSL wrapping functions.

   (e) Use the special read() and write() SSL functions to read and write to these special connected clients

   (f) If you setup your browser, you may now verify that connections to your webserver use TLSv1.0; inspect the ciphers, message authentication hash scheme, and key exchange methods used by your server.

7. Implement daemonization - we have provided you a sample of daemonizing code.

8. Submission is the same as Checkpoint 1. Tag and upload your repo in a tarball to the corresponding lab on Autolab.

5 Testing

Code quality is of particular importance for server robustness in the presence of client errors and malicious attacks. Thus, a large part of this assignment (and programming in general) is knowing how to test and debug your work. There are many ways to do this; be creative. We would like to know how you tested your server and how you convinced yourself it actually works. To this end, you should submit your test code along with brief documentation describing what you did to test that your server works. The test cases should include both generic ones that check the server functionality and those that test particular corner cases. If your server fails on some tests and you do not have time to fix it, this should also be documented (we would rather know that you are aware of the limitations of your server than think you missed a serious flaw). Several paragraphs (or even a bulleted list of things done and why) should suffice for the test case documentation.

We will be providing test scripts for each checkpoint and also the final finished server. Note however that grading will be based on additional tests that will not be provided to you. This handout lists what functions and properties of your project will be tested.

Here are some simple starting points for scripting your own external tests:

netcat

You may use netcat to send arbitrary files to your server and receive responses. Use regular bash redirection (< and >) along with ncat to achieve this.

Read `man ncat` for more information.

expect

Quoting from the expect man page:

Expect is a program that "talks" to other interactive programs according to a script. Following the script, Expect knows what can be expected from a program and what the correct response should be. An interpreted language provides branching and high-level control structures to direct the dialogue.

Python socket

This is a very simple and easy to use Python module for creating and interacting with sockets. We have used this in the first checkpoint testing script provided to you for testing your implementation of an echo server. This will be used for creating future testing programs which we will release leading the schedule for deadlines.

You can read about this module here: `http://docs.python.org/library/socket.html`.

In addition, for testing HTTP, there is a urllib2 library in Python. You can also use the requests library to create requests for your server.

6 Hand-In

Handing in code for checkpoints and the final submission deadline will be done through Autolab (`https://autolab.andrew.cmu.edu`). You are supposed to upload the tarball file for each checkpoint into the corresponding assessment on our course page of Autolab website.

6.1 Work with git

You are supposed to create your git repo on your local machine or on a **private** git repo hosted online as part of Checkpoint 1. Every checkpoint will be a git tag in this repo. To create a tag, run

```
git tag -a checkpoint-<num> -m <message> [<commit hash>]
```

with appropriate checkpoint number and custom message filled in. (Put whatever you like for the message — git won't let you omit it.) The optional commit hash can be used to specify a particular commit for the tag; if you omit it, the current commit is used. If you choose to clone your repository onto your local machine for development, be sure to use `git push --tags` to sync your work back to git server; the standard `git push` doesn't send tags.

6.2 Upload your code

To submit your code, make a tarball file of you repo after you tag it. Then login to autolab website, choose `15-441: Computer Networks (s19) -> project1cp<N>` and then upload your tarball. The grader should be finished in less than a minute but may take longer depending on system load. When it is done, your score will be shown. Only the latest score will be used.

Untarring the tarball should give us a directory named `15-441-project-1` which contains a valid git repo with tags. Your repo should contain the minimum following files:

- **Makefile** – Make sure all the variables and paths are set correctly such that your program compiles in the handin directory—not just a local machine or account. The Makefile should, by default, always build an executable named lisod.
- **All of your source code** – (files ending with .c, .h, etc. only, no .o files and no executables)
- **readme.txt** – File containing a brief description of your source tree organization.
- **tests.txt** – File containing documentation of your test cases and any known issues you have.

Late submissions will be handled according to the policy given in the course syllabus.

7 Grading

We will be providing some of the scripts but will also be running **additional tests**. Half of the points in each category (HTTP 1.1, HTTPS via TLS etc) will be based on the provided test scripts and the other half will be based on additional scripts that will be run.

- **Server core networking:** 20 points

  The grade in this section is intended to reflect your ability to write the "core" networking code. This is the stuff that deals with setting up connections, reading/writing from/to them (see the resources section below). Even if your server does not implement HTTP 1.1 etc., your project submission can get up to 20 points here. It is better to have partial functionality working solidly than lots of code that doesn't actually do anything correctly.

  Have a working `select()`-based foundation, and receive full credit here.

- **HTTP 1.1:** 20 points

  The grade in this section reflects how well you read, interpreted, and implemented HTTP 1.1. We will test all the requests specified in Section 2: HEAD, GET and POST. All requests sent to your server for this part of the testing will be valid. So a server that completely and correctly implements the specified commands, even if it does not check for invalid messages, will receive 20 points here.

  We will extensively test correct behavior for HEAD, GET, POST, and persistent connection handling. Feel free to check things via web browsers at this point.

- **HTTPS via TLS:** 10 Points

  The grade in this section reflects how well you read, interpreted, and implemented the TLS protocol for HTTP.

  Point a web browser at your server: https://xxx.xxx.xxx.xxx and verify correct connection. Obviously you will not have Certificate Authority (CA) signed certificates, but this stage should work with any web browser provided you acknowledge the security warnings and ignore them. In addition, the standard requests HEAD, GET, POST, and persistent connections should all work as before for the HTTP 1.1 implementation.

- **CGI:** 15 points

  The grade in this section reflects how well you read, interpreted, and implemented the CGI interface. This will be tested via a small Python script. The specifics for CGI can be found in RFC 3875. For the project we expect your web server to be able to run python scripts to generate dynamic pages for users. We have provided a sample script showing some basic output - cgi_script.py. You will also be creating your own cgi script that uses the following ascii art api - *http://artii.herokuapp.com/make?text=ASCII+art*. Your CGI script should take in the query parameter text along with its value, make a request to the herokuapp, and return the response to the user.

  For example the request *https://<myServer>/cgi/?text=Networks+are+cool!!!* Should return the following ascii art.

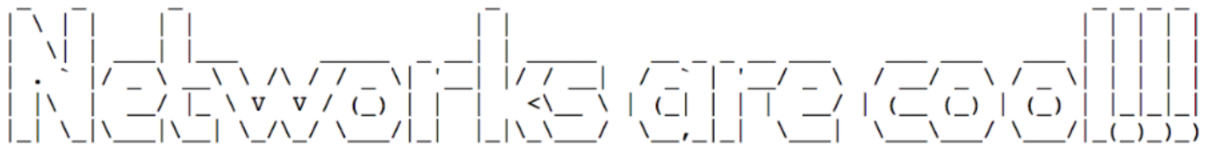

- **Robustness and Performance:** 15 points
  - Server robustness: 10 points
  - Performance: 5 points

  Since code quality is of a high priority in server programming, we will test your program in a variety of ways using a series of test cases. For example, we will send your server very long messages to test if there is a buffer overflow. We will make sure that your server does something reasonable when given an unknown request or a request with invalid headers. We will verify that your server correctly handles clients that leave abruptly without sending the proper "close" header line in HTTP 1.1. We will test that your server correctly handles concurrent requests from multiple clients without blocking inappropriately. The only exception is that your server may block while doing DNS lookups, reads from the file system, or during the execution of CGI programs.

  We will have tools that replay HTTP 1.1, TLS, and CGI interactions with your application. Note that there are many corner cases that the RFC does not specify. You will find that this is very common in "real world" programming since it is difficult to foresee all the problems that might arise. Therefore, we will not require your server pass all of the test cases in order to get a full credit on any part of the assignment. Autolab system will tell you about the errors though.

  Optimizing the performance of your web server is not a goal of this project. However, we will penalize you for egregious design decisions that will fundamentally limit performance. Examples of poor design decisions include very inefficient parsing (e.g., not using lex and yacc but inefficient custom parsing), inefficient buffer management (e.g., unnecessary calls to malloc), or poor choice of data structures (e.g., linked lists instead of hash tables).

- **Design, Style and Your Tests: 20 points**
  - Design: 5 points
  - Style: 5 points
  - Your Tests: 10 points

  Poor design, documentation, or code structure will probably reduce your grade by making it hard for you to produce a working program and hard for the grader to understand it; egregious failures in these areas will cause your grade to be lowered even if your implementation performs adequately.

  Document code using Doxygen-style comments.

  We expect you to come up with test scripts and there are 10 points for test scripts/manual testing. We expect you to provide a list of tests in test.txt with a brief description of what each test is testing. Quantity of test scripts is not a criteria. We expect your test scripts to test different parts of your system thoroughly. Also, verifying the correctness of tests requires logging, so we expect you to implement a logging system (see Section 4.2).

8 Getting Started

This section gives suggestions for how to approach the project. Naturally, other approaches are possible, and you are free to use them.

- **Start early!** The hardest part of getting started tends to be getting started. Remember the 90-90 rule: the first 90% of the job takes 90% of the time; the remaining 10% takes the other 90% of the time. Starting early gives you time to ask questions. For clarifications on this assignment, post to Piazza and read project updates on the course web page. Talk to your classmates. While you need to write your own original program, we expect conversation with other people facing the same challenges to be very useful. Come to office hours. The course staff is here to help you.
- Read the RFCs selectively. RFCs are written in a style that you may find unfamiliar. However, it is wise for you to become familiar with it, as it is similar to the styles of many standards organizations. We don't expect you to read every page of the RFC, especially since you are only implementing a small subset of the full protocol, but you may well need to re-read critical sections a few times for the meaning to sink in.
- Begin by taking a cursory first pass over the RFCs. Do not focus on the details; just try to get a sense of how they work at a high level. Understand the role of the server. Understand what error conditions are possible, and how they are used. You may want to print the RFCs, and mark them up to indicate which parts are important for this project, and which parts are not needed. You may need to re-read these sections several times.
- Next, take a second pass over the RFCs. You will want to read all of them together. Again, do not focus on the details; just try to understand the requests and responses

at a high level. As before, you may want to mark up a printed copy to indicate which parts of the RFCs are important for the project, and which parts are not needed.

- Now, go back and read with an eye toward implementation. Mark the parts which contain details that you will need to write your server. Start thinking about the data structures (input and output buffers, etc.) your server will need to maintain. What information needs to be stored about each client while servicing requests (maybe an HTTP 1.1 finite state machine per client, etc.)?

- Get started with a simple server that accepts connections from multiple clients. It should take any message sent by any client, and "echo" that message back to its sender. This server will not be compatible with HTTP clients, but the code you write for it will be useful for your final server. Writing this simpler server will let you focus on the socket programming aspects of a server without worrying about the details of the protocols. Test this simple server with the provided Python script in Checkpoint 1.

- Next, enhance the starter code we have provided for HTTP parsing. Apply all the RFC knowledge you have gathered from previous steps and try to convert them into rules. After you combine this parser along with the echo server, you should be ready for Checkpoint 1.

- At this point, you are ready to write a standalone HTTP 1.1 server. But do not try to write the whole server at once. Decompose the problem so that each piece is manageable and testable. For each request, identify the different cases that your server needs to handle. Find common tasks among different commands and group them into procedures to avoid writing the same code twice. You might start by implementing the routines that read and parse commands. Then implement commands one by one, testing each with `telnet`.

- Thoroughly test your server. Use the provided scripts to test basic functionality. For further testing, use `telnet`, a web browser, or replay scripts. Learn Python from our scripts and as we go to make repeatable "regression tests"—every time you implement a new feature you use regression tests to see if anything broke.

- Make sure to check the return code of all system calls and handle errors appropriately. Temporary failures (e.g., EINTR) should not cause your server to abort or exit in failure. Fatal errors can be dealt with via a perror() call and exiting—but try to clean up open file descriptors and sockets nicely even when fatally exiting.

- Be liberal in what you accept and conservative in what you send [10]. Following this guiding principle of Internet design will help ensure your server works with many different and unexpected client behaviors.

- Code quality is important. Make your code modular and extensible where possible. You should probably invest an equal amount of time in testing and debugging as you do writing. Also, debug incrementally. Write in small pieces and make sure they work before going on to the next piece. Your code should be readable and commented. Not only should your code be modular, extensible, readable, etc, most importantly, it should be your own!

- You may want to consider turning warnings into errors to avoid bad programming style. Do this by passing `-Werror` to `gcc` during compilation.
- If you have a question about a project handout or a technical issue, there is an excellent chance that other students have the same question. Please read Piazza to see if there has been traffic and consider posting your questions there.

9 Resources

For information on network programming, the following may be helpful:

- Beej's Guide [8]
- Class Textbook – Sockets, etc
- Class Piazza – Announcements, clarifications, etc
- Class Website – Announcements, errata, etc
- Computer Systems: A Programmer's Perspective (CS 15-213 text book)[15]
- BSD Sockets: A Quick And Dirty Primer[16]
- An Introductory 4.4 BSD Interprocess Communication Tutorial[17]
- Unix Socket FAQ[18]
- Sockets section of the GNU C Library manual
  - Installed locally: info libc
  - Available online: GNU C Library manual[19]
- man pages
  - Installed locally (e.g. man socket)
  - Available online: the Single Unix Specification[20]
- Other Google Groups / Stackoverflow - Answers to almost anything[21]

References

[1] RFC 2396: `http://www.ietf.org/rfc/rfc2396.txt`

[2] RFC 2616: `http://www.ietf.org/rfc/rfc2616.txt`

[3] RFC 2818: `http://www.ietf.org/rfc/rfc2818.txt`

[4] RFC 3875: `http://www.ietf.org/rfc/rfc3875`

[5] Apache: `http://httpd.apache.org/`

[6] Wireshark: `http://www.wireshark.org/`

[7] Open SSL: `https://www.openssl.org/docs/manmaster/man3/`

[8] Beej's Guide: http://beej.us/guide/bgnet/output/html/singlepage/bgnet.html

[9] GNU Make Manual: http://www.gnu.org/software/make/manual/make.html

[10] RFC 1122 http://www.ietf.org/rfc/rfc1122.txt, page 11

[11] ElectricFence: http://perens.com/FreeSoftware/ElectricFence/

[12] Common Log Format: https://httpd.apache.org/docs/1.3/logs.html#common

[13] No-IP: https://www.noip.com/free

[14] Valgrind: http://valgrind.org/

[15] CSAPP: http://csapp.cs.cmu.edu

[16] http://www.cis.temple.edu/~ingargio/old/cis307s96/readings/docs/sockets.html

[17] http://docs.freebsd.org/44doc/psd/20.ipctut/paper.pdf

[18] http://www.developerweb.net/forum/forumdisplay.php?s=f47b63594e6b831233c4b8ebaf10a614&f=70

[19] http://www.gnu.org/software/libc/manual/

[20] http://www.opengroup.org/onlinepubs/007908799/

[21] http://groups.google.com