

CSc 450/550: Computer Communications and Networks (Summer 2007)

Lab Project 1: A Simple Multi-Thread Web Server

Spec Out: May 11, 2007

Demo Due: May 30, 2007

Code Due: June 1, 2007

1 Introduction

In this project, students will build a simple, multi-thread “Web” server in C or C++. The project will allow students to refresh their C or C++, `socket` and `pthread` programming skills, and to better understand the Hyper-Text Transfer Protocol (HTTP) underneath the World-Wide Web.

2 Background

2.1 `socket` API

`socket` is the Application Programming Interface (API) to network services in many operating systems. For a “server”-like application, normally you will need to use the following system calls.

- `socket()`: to create a new socket
- `bind()`: to associate the socket to a local address
- `listen()`: to wait for an incoming connection request
- `accept()`: to accept the incoming connection request with a new socket
- `recv()`: to read from the socket
- `send()`: to write to the socket
- `close()`: to close the socket and release the resources allocated

You will want to read the manual page (e.g., `man socket`) to better understand these system calls and their input arguments and return values. You may need to use other `socket`-related system calls or auxiliary library calls as well. It is very important to check the return value of these function calls in your program to determine whether the intended operation is successful.

For more information on `socket` programming, please see [1].

27 2.2 pthread Library

28 pthread is the POSIX standards-based API to multi-threading programming in many operating
29 systems. The following library calls are used to create, terminate and synchronize threads.

- 30 • `pthread_create()`: to create a new thread
- 31 • `pthread_exit()`: to terminate the calling thread
- 32 • `pthread_join()`: to suspend the calling thread and to wait for the target thread to finish

33 All threads in the same process share most data. To provide finer thread synchronization, you
34 may want to use mutual exclusion (`pthread_mutex_lock()` and `pthread_mutex_unlock()`) and
35 condition variables (`pthread_cond_wait()` and `pthread_cond_signal()`) in some cases.

36 Again, you may want to read the `pthread` manual page (e.g., `man pthread_create`).

37 For more information on `pthread` programming, please see [2].

38 2.3 HTTP Protocol

39 HTTP is a client-server, request-response application-layer protocol for the Web. In this project,
40 only a very simplified version of HTTP/1.0 is to be implemented.

41 2.3.1 Simplified HTTP Request

42 As being stated in the HTTP/1.0 specification: “A request message from a client to a server includes,
43 within the first line of that message, the method to be applied to the resource, the identifier of the
44 resource, and the protocol version in use.” [3]

45 For example, when `wget -d www.csc.uvic.ca`, it shows

46 Message	Explanation (required or ignored by the simple web server)
<code>GET / HTTP/1.0</code>	method, request URI, HTTP version (required)
<code>User-Agent: Wget/1.10.2</code>	Request header (ignored)
47 <code>Accept: */*</code>	Request header (ignored)
<code>Host: www.csc.uvic.ca</code>	Request header (ignored)
<code>Connection: Keep-Alive</code>	Request header (ignored)
	a blank line indicating the end of request headers (required)

48 The simple web server will read the request line, which includes **method**, **request URI** and
49 **HTTP version** separated by **string delimiters** (e.g., blank spaces or tabs), but will ignore all
50 request headers, until it reaches the end of request headers indicated by **a blank line**.

51 **Method** and **HTTP version** are not case sensitive; however, **URI** is case sensitive.

52 The simple web server only supports the `GET` method, which obtains an HTTP object (often an
53 HTML file) from the web server.

54 **URI** identifies the HTTP object. For example, `GET / HTTP/1.0` tries to obtain the default file
55 `index.html` from the root of the web server document directory. If the web server’s root document
56 directory is `/tmp/www`, then `/tmp/www/index.html`, if existent, is retrieved and returned by the
57 web server; otherwise, a Not Found error message is returned instead. `GET /icons/new.gif` will
58 let the web server retrieve `/tmp/www/icon/new.gif`.

59 The simple web server only recognizes HTTP/1.0 as the valid HTTP version supported.

60 The web server will return a **Bad Request** error message if an invalid request line or an incom-
61 plete request message is encountered. The minimal, valid request message contains a valid request
62 line and a blank line indicating the end of the request.

63 2.3.2 Simplified HTTP Response

64 As being stated in the HTTP/1.0 specification: “After receiving and interpreting a request message,
65 a server responds in the form of an HTTP response message.” [3]

66 When `wget -d www.csc.uvic.ca`, the web server returns

67

Message	Explanation (required or omitted in the simple web server)
HTTP/1.1 200 OK	HTTP version, status code, reason phrase (required)
Date: Wed, 02 May ...	response header (omitted)
Server: Apache/2. ...	response header (omitted)
68 X-Powered-By: PHP/4 ...	response header (omitted)
Connection: c ...	response header (omitted)
Content-Type: text/ ...	response header (omitted)
	a blank line indicating the end of response headers (required)
index.html content	HTTP object(s) returned (required, if any)

69 The simple web server will return the response line, which includes **HTTP version, status**
70 **code** and **reason phrase** separated by **blank spaces**, but will omit all response headers shown
71 above. However, the simple web server will return **a blank line** indicating the end of response
72 headers and the start of the returned HTTP object, if any.

73 The response line is not case sensitive, but it is suggested to follow the convention shown above.

74 The simple web server only returns HTTP/1.0 as the valid HTTP version supported.

75 The simple web server only supports the following list of status codes and reason phrases.

76

Status Code	Reason Phrase	Explanation
200	OK	good request with the requested object to be returned
400	Bad Request	bad request not understood by the server
404	Not Found	good request with no matching request object

77 If the simple web server can understand the request, and the request object is successfully
78 retrieved, it will return HTTP/1.0 200 OK followed by a blank line and the content of the requested
79 object. If the simple web server cannot understand the request, it will return HTTP/1.0 400 **Bad**
80 **Request** followed by a blank line. If the simple web server can understand the request, but the
81 requested object is not available (e.g., nonexistence, bad file permission, etc), it will return HTTP/1.0
82 404 **Not Found** followed by a blank line.

83 The simple web server does not have to support persistent connections.

84 For more information on HTTP/1.0, please see [3].

85 3 Design

86 Students will have the freedom of designing their simple, multi-thread web server to be implemented
87 in C or C++ with `socket` API and `pthread` library. The following is just a possible design.

88 When the server is invoked, among other things, it will create a “dispatcher” thread to accept
89 incoming TCP connections, which transport the HTTP request and response. It can also create a
90 few “worker” threads in advance to be assigned by the dispatcher to incoming connections, or these
91 threads are created on demand by the dispatcher when a connection comes. The worker thread
92 actually reads and processes the HTTP request, and returns the HTTP response and the requested
93 object if the request is successful (or an error message otherwise). The worker threads could be
94 terminated after an HTTP request is served, or be kept in an idle thread pool for future requests.

95 4 Requirements

96 4.1 Basic Features

97 Basic features are required in all implementations in order to get the full marks for this lab project.

98 4.1.1 Invoking the Server

99 The syntax to run the simple web server is

```
100 ./sws <port> <directory>
```

101 This will invoke the simple web server binary `sws` in the current directory, and instruct the web
102 server to listen at the TCP `port` for incoming connections and to retrieve requested objects under
103 `directory`. If a wrong syntax is used when invoking the server, the server should print out error
104 messages showing the proper usage and exit gracefully.

105 When it is invoked successfully, for example, the simple web server will print out the following
106 message if being invoked as `./sws 8080 /tmp/www`

```
107 sws is running on port 8080 and serving /tmp/www  
108 press 'q' to quit ...
```

109 The client is not allowed to use `../..` to retrieve objects out of the root document `directory`;
110 such requests will be responded with HTTP/1.0 400 Bad Request by the server.

111 4.1.2 Server Operations

112 When running, the simple web server should respond to incoming requests at TCP `port` for all
113 network interfaces on the machine running the server.

114 Once a request is served, the server should print out a log message in the following format.

```
115 MMM DD HH:MM:SS Client-IP:Client-Port request-line response-line [filename]
```

116 For example, if the server, on the noon of May 11, successfully served a request `GET / HTTP/1.0`
117 from the client at port 4096 on host `192.168.0.100`, it should print out

```
118 May 11 12:00:00 192.168.0.100:4096 GET / HTTP/1.0 HTTP/1.0 200 OK /tmp/www/index.html
```

119 `filename` indicates the location of the HTTP object returned, if the request is successfully
120 served. The log messages should be ordered chronically and not be scrambled due to the fact that
121 there are multiple worker threads.

122 4.1.3 Multi-Threading

123 The simple web server should use multi-threading (up to five “worker” threads) to handle multiple,
124 concurrent HTTP requests in a non-blocking fashion. When there are more than five concurrent
125 connection requests, the additional requests (at least five extra requests) should be queued in
126 `listen()`. In other words, the simple web server should be able to handle ten concurrent HTTP
127 requests without dropping any of them.

128 4.1.4 Terminating the Server

129 The simple web server should continue serving HTTP requests until the `q` key is pressed in the
130 terminal running the server. When the `q` key is pressed, the server should finish serving all accepted
131 requests, close all created sockets, and release all allocated resources before exiting gracefully.

132 After terminating the server, or if the server is aborted, some socket resources may still be used
133 by the system. If you rerun the server on the same port immediately (or if there is already a
134 program on the same port), normally you will get a `bind()` error. You can either terminate the
135 other program, wait for a while, use another port, or use `setsockopt()` with level `SOL_SOCKET` and
136 option `REUSEADDR` before `bind()` to allow port reuse.

137 4.2 Bonus Features

138 Only a very simplified version of HTTP/1.0 is to be supported by the multi-thread web server.
139 However, students have the option to extend their design and implementation to include more
140 features in the HTTP protocol (e.g., persistent HTTP connection, HTTP request pipelining, etc)
141 and to improve server performance (e.g., caching, etc).

142 If you want to design and implement a bonus feature, you should contact the course instructor
143 for permission at least one week before the code due date and clearly indicate the feature during
144 project demo and in code submission. The credit for correctly implemented bonus features will not
145 exceed 20% of the full marks for this lab project.

146 5 Self-Testing

147 You can test your multi-thread web server with any HTTP/1.0-compliant client, even Telnet. Using
148 Telnet, you can connect to your web server running at `port` on `host` by

```
149 telnet <host> <port>
```

150 You can then manually type the HTTP request after Telnet says

151 Escape character is '^]'.
152

153 And the HTTP response should follow after a correct HTTP request is successfully processed.
154 You can also use multiple Telnet sessions to test the multi-threading feature of your web server.

155 To assist your testing, you will be provided a gzipped sample document directory file `www.tar.gz`,
156 and you can recover the sample document directory in `/tmp` by

```
156 mv www.tar.gz /tmp  
157 cd /tmp  
158 tar -zxvf www.tar.gz
```

159 A sample request sequence `http.request` is also included in `www.tar.gz`
160 However, your simple web server might be tested against different document directories and
161 request sequences during project demo and code inspection.

162 6 Demonstration

163 Your simple web server should be demonstrated in the lab section for which you have registered on
164 the demo due date. **Lab projects not demonstrated will have their code submission not**
165 **marked. It is expected that your lab project be working at the time of demonstration.**
166 During the project demo, lab instructors will go through a demo checklist together with the student,
167 and then provide the checklist to the student. Students will have the chance to improve their design
168 and implementation until the code due date.

169 The demo is intended to allow students to demonstrate their projects and to help them improve
170 their design and implementation, not to be coding or debugging assistance, and only focuses on
171 required features (and bonus features if indicated). There is no guarantee on the correctness and
172 grade of the project, which can only be determined after the code inspection.

173 7 Submission

174 The entire lab project, including the code and documentation, should be submitted electrically
175 through `csc4501` (1 is for letter L) at `http://www.csc.uvic.ca/~submit/index.cgi` on or before
176 the code due date.

177 Only the source code (including header files and Makefile) and documentation (including Readme)
178 should be included in a single `tar.gz` file to be submitted. No object or binary files are included
179 in the submission. If `directory` is your project directory, to create such a gzipped tarball, you can

```
180 cd direcoty  
181 tar -zcvf p1.tar.gz .
```

182 This packing and naming convention should be strictly followed to allow your submission to be
183 properly located for grading.

184 In `directory`, you need to include a `Makefile`, which compiles and builds the final binary
185 executable (`sws`) automatically by typing

```
186 make
```

187 The same `Makefile` also removes all object and binary files when you type in

```
188 make clean
```

189 All projects will be tested on `linux.csc.uvic.ca`

190 In `directory`, you also need to include a `Readme` plain text file, which contains your student
191 number, registered lab section and a brief description of your design and code structure, as well as
192 allowed bonus features, if any.

193 The code itself should be sufficiently self-documented. For more information on acceptable
194 coding style, please see [4].

195 **IMPORTANT:** All submitted work should be yours. If you have used anything out there,
196 even a small component in your implementation, you should credit and reference properly, and
197 your contribution can be determined accordingly. For academic integrity policies, please see [5].

198 8 Marking

199 This lab project is worth 15% in the final grade of this course for CSc 450 students, and 10% for
200 CSc 550 students.

201 For mark posting and appeal policies, please see the official course outline at [5].

202 References

203 [1] <http://beej.us/guide/bgnet/>

204 [2] <http://www.llnl.gov/computing/tutorials/pthreads/>

205 [3] <http://www.w3.org/Protocols/rfc1945/rfc1945>

206 [4] <http://www.csc.uvic.ca/~csc4501/references/Code-Style.html>

207 [5] <http://courses.seng.engr.uvic.ca/courses/2007/summer/csc/450>