


CSE 333

Section 8

Booooooooooost & HTTP



Logistics

Friday:

Exercise 16 @ 10:30 am

Monday:

Exercise 17 @ 10:30 am

Thursday:

HW4 @ 11:59 pm

BOOOOOOOOST

BOOST

Boost is a free C++ library that provides support for various tasks in C++

- **Note:** Boost does NOT follow the Google style guide!!!

Boost adds many string algorithms that you may have seen in Java

- Include with `#include <boost/algorithm/string.hpp>`

We are showcasing a few we think could be useful for HW4, but more can be found here:

- https://www.boost.org/doc/libs/1_60_0/doc/html/string_algo.html

trim

```
void boost::trim(string& input);
```

- Removes all leading and trailing whitespace from the string
- `input` is an input *and* output parameter (non-const reference)

```
string s("    HI    ");  
boost::algorithm::trim(s);
```

```
// results in s == "HI"
```

replace_all

```
void boost::replace_all(string& input, const string& search,  
                           const string& format);
```

- Replaces all instances of `search` inside `input` with `format`

```
string s("ynrnrt");  
boost::algorithm::replace_all(s, "nr", "e");  
  
// results in s == "yeet"
```

replace_all

```
void boost::replace_all(string& input, const string& search,  
                           const string& format);
```

- Replaces all instances of `search` inside `input` with `format`

```
string s("queue?");  
boost::algorithm::replace_all(s, "que", "q");  
// results in s == "qe?"
```

`replace_all()` guarantees that ‘format’ will be in the final result if-and-only-if ‘search’ existed.

`replace_all()` makes a *single* pass over input.

split

```
void boost::split(vector<string>& output,  
    const string& input,  
    boost::PredicateT match_on,  
    boost::token_compress_mode_type compress);
```

- Split the string by the characters in `match_on`

```
boost::PredicateT boost::is_any_of(const string& tokens);
```

- Returns predicate that matches on any of the characters in `tokens`

split Examples

```
vector<string> tokens;  
string s("I-am--split");
```

```
boost::split(tokens, s, boost::is_any_of("-"),  
              boost::token_compress_on);  
// results in tokens == ["I", "am", "split"]
```

```
boost::split(tokens, s, boost::is_any_of("-"),  
              boost::token_compress_off);  
// results in tokens == ["I", "am", "", "split"]
```

Exercise 1

Write a function that takes in a string that contains words separated by whitespace and returns a vector that contains all of the words in that string, in the same order as they show up, but with no duplicates. Ignore all leading and trailing whitespace in the input string.

Example:

`RemoveDuplicates(" Hi I'm sorry jon sorry hi hihi hi hi ")`
should return the vector `["Hi", "I'm", "sorry", "jon", "hi", "hihi"]`

```
vector<string> RemoveDuplicates(const string& input){
    string copy(input);
    boost::algorithm::trim(copy);
    std::vector<string> components;
    boost::split(components, copy, boost::is_any_of(" \t\n"),
                  boost::token_compress_on);
    std::vector<string> result;
    for (uint i = 0; i < components.size(); ++i) {
        bool unique = true;
        for (uint j = 0; j < i && unique; ++j) {
            unique &= !(components[i] == components[j]);
        }
        if (unique) {
            result.push_back(components[i]);
        }
    }
    return result;
}
```

HTTP Review

HTTP Review

1. What does HTTP stand for?

Hyper**T**ext **T**ransfer **P**rotocol

2. What layer does HTTP reside in?

Application **L**ayer

HTTP Request Format

[METHOD] [request-uri] HTTP/[version]\r\n

[headerfield1]: [fieldvalue1]\r\n

[headerfield2]: [fieldvalue2]\r\n




[...]

[headerfieldN]: [fieldvalueN]\r\n

\r\n

[request body, if any]

HTTP Methods

| | | |
|--|----------------|---|
|  | GET | The GET method requests a representation of the specified resource. Requests using GET should only retrieve data. |
|  | HEAD | The HEAD method asks for a response identical to that of a GET request, but without the response body. |
|  | POST | The POST method is used to submit an entity to the specified resource, often causing a change in state or side effects on the server. |
| | PUT | The PUT method replaces all current representations of the target resource with the request payload. |
| | DELETE | The DELETE method deletes the specified resource. |
| | CONNECT | The CONNECT method establishes a tunnel to the server identified by the target resource. |
| | OPTIONS | The OPTIONS method is used to describe the communication options for the target resource. |
| | TRACE | The TRACE method performs a message loop-back test along the path to the target resource. |
| | PATCH | The PATCH method is used to apply partial modifications to a resource. |

HTTP Response Format

HTTP/[version] [status code] [reason]\r\n

[headerfield1]: [fieldvalue1]\r\n

[headerfield2]: [fieldvalue2]\r\n

[...]

[headerfieldN]: [fieldvalueN]\r\n

\r\n

[response body, if any]

HTTP Response Status Codes

- HTTP/1.1 200 OK
 - The request succeeded and the requested object is sent
- HTTP/1.1 404 Not Found
 - The requested object was not found
- HTTP/1.1 301 Moved Permanently
 - The object exists, but its name has changed
 - The new URL is given as the “Location: ” header value
- HTTP/1.1 500 Server Error
 - The server had some kind of unexpected error

Version

```
HTTP/1.1 200 OK
Date: Mon, 21 May 2018 07:58:46 GMT
Server: Apache/2.2.32 (Unix) mod_ssl/2.2.32 OpenSSL/1.0.1e-fips
mod_pubcookie/3.3.4a mod_uwa/3.2.1 Phusion_Passenger/3.0.11
Last-Modified: Mon, 21 May 2018 07:58:05 GMT
ETag: "2299e1ef-52-56cb2a9615625"
Accept-Ranges: bytes
Content-Length: 82
Vary: Accept-Encoding,User-Agent
Connection: close
Content-Type: text/html
Set-Cookie:
bbbbbbbbbbbbbbbb=DBMLFDMJCGAOILMBPIIAAIFLGBAKOJNNMCJIKKBKCDMDEJHMPONHCILPIBL
ADEAKCIABMEEPAOPMMKAOLHOKJMIGMIDKIHNCANAPHMFMBLBABPFENPDANJAPIBOIOOOD;
HttpOnly

<html><body>
<font color="chartreuse" size="18pt">Awesome!!</font>
</body></html>
```

Status

Headers

Request body

HTTP REQUEST DEMO (telnet)

Using Telnet with HW4

1. Launch the server

```
./http333d <port> ../projdoks/ unit_test_indices/*
```

2. Connect with telnet

```
telnet <HostName> <port>
```

3. Write an HTTP request and send it

4. To exit telnet:

- **Ctrl+]** then **Ctrl+d**

Writing an HTTP Request

- Example HTTP Request layout can be found in `HttpRequest.h`
- Example file request:
 - `GET /static/test_tree/books/artofwar.txt HTTP/1.1`
- Example query request:
 - `GET /query?terms=books+of+war HTTP/1.1`
- To send a request, **hit [Enter] twice**
- Compare the output of `solution_binaries/http3d` to `./http3d`

Exercise 4

Exercise 4

Write a function called `ExtractRequestLine` that takes in a well-formatted HTTP request as a `string` and returns a `map` with the keys as `method`, `uri`, `version` and the values from the corresponding request.

Example Input:

```
"GET /index.html HTTP/1.1\r\nHost: www.mywebsite.com\r\nConnection: keep-alive\r\nUpgrade-Insecure-Requests: 1\r\n\r\n"
```

Map Returned:

```
{  
  "method"   : "GET"  
  "uri"      : "/index.html"  
  "version"  : "HTTP/1.1"  
}
```

Exercise 4

```
map<string, string> ExtractRequestLine(const string& request) {  
    vector<string> lines;  
    boost::split(lines, request, boost::is_any_of("\r\n"),  
                 boost::token_compress_on);  
    vector<string> components;  
    string firstLine = lines[0];  
    boost::split(components, firstLine, boost::is_any_of(" "),  
                 boost::token_compress_on);  
    map<string, string> res;  
    res["method"] = components[0];  
    res["uri"] = components[1];  
    res["version"] = components[2];  
    return res;  
}
```

Threads

“Computers are really dumb. They can only do a few things like shuffling around numbers, but they do them really really fast so that they appear smart.”

-Hal Perkins

Threads are just a way of making computers appear to do multitasking, regardless of whether they are running one or more CPUs

Threads

- Contained within a process.
- Multiple threads can exist within the same process.
 - Every process starts with one thread of execution, but it can spawn more.
- Threads in a single process share one address space
 - Instructions (code)
 - Static (global) data
 - Dynamic (heap) data
 - Environment variables, open files, sockets, etc.
- Easy communication (put something in shared memory)
- Synchronization often uses locks (like mutexes)

POSIX threads (pthreads)

- The POSIX standard provides APIs for creating and manipulating threads.
- Part of the standard C/C++ libraries, declared in `pthread.h`
- Core pthread functions:
 - `pthread_create` - “Go do this {function}”
 - `pthread_exit` - “I’m done with my task!”
 - `pthread_join` - “I’ll wait for you to report back your result”
 - `pthread_cancel` - “I changed my mind, you can stop now”
 - `pthread_detach` - “You’re free now, go forth and prosper”

pthread_create

```
#include <pthread.h>
int pthread_create( pthread_t *thread,
                   const pthread_attr_t *attr,
                   void *(*start_routine) (void *),
                   void *arg );
```

- `pthread_create` creates a new thread and calls `start_routine` with `arg` as its parameter.
- `pthread_create` arguments:
 - **thread:** Pointer to a unique identifier for the new thread. (output parameter)
 - **attr:** An attribute object that may be used to set thread attributes. Use NULL for the default values.
 - **start_routine:** The C routine that the thread will execute once it is created.
 - **arg:** A single argument that may be passed to `start_routine`. It must be passed by reference as a pointer cast of type void. NULL may be used if no argument is to be passed.
- Compile and link with `-pthread`.

Threads - Quick Check

```
MyClass onTheStack;  
pthread_t child;  
pthread_create(&child, nullptr, foo, &onTheStack);
```

onTheStack is on the parent thread's stack. However, each thread has its own stack! Can we still access onTheStack from the child? Why or why not?

Yes! All threads share an address space

Terminating Threads

- There are several ways in which a thread may be terminated:
 - The thread returns normally from its starting routine; Its work is done.
 - The thread makes a call to the `pthread_exit` subroutine - whether its work is done or not.
 - The thread is canceled by another thread via the `pthread_cancel` routine.
 - The entire process is terminated due to making a call to either the `exec()` or `exit()`.
 - If `main()` finishes first, without calling `pthread_exit` explicitly itself.

pthread_exit

```
void pthread_exit(void *retval);
```

- Allows the user to terminate a thread and to specify an optional termination status parameter, *retval*.
- In subroutines that execute to completion normally, you can often dispense with calling `pthread_exit()`.
- Calling `pthread_exit()` from `main()`:
 - If `main()` finishes before the threads it spawned and does not call `pthread_exit()` explicitly, all the threads it created will terminate.
 - To allow other threads to continue execution, the main thread should terminate by calling `pthread_exit()` rather than `exit()`.

Threads - Gotchas

- Resources (heap-allocated storage, file descriptors, etc)
 - Often shared between multiple threads
 - Must be allocated / deallocated *exactly once*
 - Don't use deallocated resources from other threads

```
buf = new int[BUFSIZE];  
...  
if (!handleRequest(buf, req, len)) {  
    delete[] buf; // buf was allocated in this thread  
    close(fd); // is somebody else going to try to use fd??  
    pthread_exit(nullptr);  
}
```


pthread_join

```
int pthread_join(pthread_t thread, void **retval);
```

- Synchronization between threads.
- `pthread_join` blocks the calling thread until the specified thread terminates and then the calling thread joins the terminated thread.
- Only threads that are created as joinable can be joined; a thread created as detached can never be joined. (Refer `pthread_create`)
- The target thread's termination return status can be obtained if it was specified in the target thread's call to `pthread_exit()`.

Demo: *pthreads.cc*

Locking - mutex

```
int pthread_mutex_init(pthread_mutex_t *mutex,  
                        const pthread_mutexattr_t *attr);
```

- Initializes the mutex lock pointed to by `mutex` with lock attributes specified by `attr`.
- `Attr` can be `null`.

```
int pthread_mutex_lock(pthread_mutex_t *mutex);
```

- Grabs the lock

```
int pthread_mutex_unlock(pthread_mutex_t *mutex);
```

- Releases the lock

```
int pthread_mutex_destroy(pthread_mutex_t *mutex);
```

- Destroys the lock

Threads – Locking

- Locking is hard.
 - Too much, and performance is *worse than sequential*
 - Too little, and threads clash - *often unexpected results*
 - Not careful, and **deadlock** freezes your program forever!

```
pthread_mutex_lock(&lock);  
if (!do_computation(resource)) {  
    printf("Error doing computation\n");  
    return false; // !!!  
}  
pthread_mutex_unlock(&lock);  
return true;
```

Demo total.cc & total_locking.cc

More examples of pthreads usage

- From sequential to concurrent merge sort
 1. `merge_sort.cc`
 2. `c4_merge_sort.cc`
 3. `concurrent_merge_sort.cc`