

# Project 1: Snowcast

*Due: 11:59 PM, Sep 24, 2015*

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>Protocol</b>	<b>2</b>
2.1	Client to Server Commands . . . . .	2
2.2	Server to Client Replies . . . . .	3
2.3	Invalid Conditions . . . . .	3
2.3.1	Server . . . . .	3
2.3.2	Client . . . . .	4
<b>3</b>	<b>Implementation Requirements</b>	<b>4</b>
3.1	Correctness . . . . .	5
3.2	Clients . . . . .	5
3.2.1	UDP Client . . . . .	5
3.2.2	TCP Client . . . . .	5
3.3	Server . . . . .	6
<b>4</b>	<b>Testing</b>	<b>6</b>
4.1	Rate Monitor . . . . .	7
4.2	Reference Implementations . . . . .	7
<b>5</b>	<b>Grading</b>	<b>7</b>
5.1	Milestone - 30% . . . . .	7
5.2	Program - 65% . . . . .	8
5.3	README - 5% . . . . .	8
<b>6</b>	<b>Handin and Interactive Grading</b>	<b>8</b>
<b>A</b>	<b>Useful Hints/Tips</b>	<b>9</b>

## 1 Introduction

You will be implementing a simple Internet Radio Station. The purpose of this assignment is to become familiar with sockets and threads, and to get you used to thinking about network protocols.

If you're unfamiliar with sockets or threads, you should read the pages about them linked from the course web page. As always, e-mail us at `cs168tas@cs.brown.edu` or to come to one of our office hours if you have further questions.

## 2 Protocol

This assignment has two parts: the server, which streams songs, and a pair of clients for connecting to the server and receiving songs.

There are two kinds of data being sent between the server and the client. One is the control data. The client uses this data to specify which station to listen to and the server uses it to give the client song information. The other kind is the song data, which the server reads from song files and streams to the client. You will be using TCP for the control data and UDP for the song data.

The client and server communicate control data by sending each other **messages** over the TCP connection.

### 2.1 Client to Server Commands

The client sends the server messages called **commands**. There are two commands the client can send the server, in the following format.

Hello:

```
uint8_t commandType = 0;
uint16_t udpPort;
```

SetStation:

```
uint8_t commandType = 1;
uint16_t stationNumber;
```

A `uint8_t`<sup>1</sup> is an unsigned 8-bit integer. A `uint16_t` is an unsigned 16-bit integer. Your programs MUST use **network byte order**.<sup>2</sup> So, to send a Hello command, your client would send exactly three bytes to the server.

The Hello command is sent when the client connects to the server. It tells the server what UDP port the server should be streaming song data to.

The SetStation command is sent to pick an initial station or to change stations. `stationNumber` identifies the station.

---

<sup>1</sup>You can use these types from C if you `#include <inttypes.h>`.

<sup>2</sup>Use the functions `htons`, `htonl`, `ntohs` and `ntohl` to convert from network to host byte order and back.

## 2.2 Server to Client Replies

There are three possible messages called **replies** the server may send to the client:

Welcome:

```
uint8_t  replyType = 0;
uint16_t numStations;
```

Announce:

```
uint8_t  replyType = 1;
uint8_t  songnameSize;
char     songname[songnameSize];
```

InvalidCommand:

```
uint8_t  replyType = 2;
uint8_t  replyStringSize;
char     replyString[replyStringSize];
```

A **Welcome** reply is sent in response to a **Hello** command. Stations are numbered sequentially from 0, so a **numStations** of 30 means 0 through 29 are valid. A **Hello** command, followed by a **Welcome** reply, is called a **handshake**.

An **Announce** reply is sent on two occasions: after a client sends a **SetStation** command, or when the station a client is listening to changes its song. **songnameSize** represents the length, in bytes, of the filename, while **songname** contains the filename itself. The string must be formatted in ASCII and must **not be null-terminated**. So, to announce a song called **Gimme Shelter**, your client would send the **replyType** byte, followed by a byte whose value is 13, followed by the 13 bytes whose values are the ASCII character values of **Gimme Shelter**.

## 2.3 Invalid Conditions

Since neither the client nor the server may assume that the program with which it is communicating is compliant with this specification, they **MUST** both be able to behave correctly when the protocol is used incorrectly.

### 2.3.1 Server

On the server side, an **InvalidCommand** reply is sent in response to any invalid command. **replyString** should contain a brief error message explaining what went wrong. Give helpful strings stating the reason for failure. If a **SetStation** command was sent with 1729 as the **stationNumber**, a bad **replyString** is “Error, closing connection.”, while a good one is “Station 1729 does not exist.”. To simplify the protocol, whenever the server receives an invalid command, it **MUST** reply with an **InvalidCommand** and then **close the connection** to the client that sent it.

Invalid commands happen in the following situations:

- **SetStation**
  - The station given does not exist.

- The command was sent before a `Hello` command was sent. The client must send a `Hello` command before sending any other commands.
- If the command was sent before the server responded to a previous `SetStation` by sending an `Announce` reply, then your server MAY reply to this with an `InvalidCommand`. This means that your client should be careful and wait for an `Announce` before sending another `SetStation`, but your server can be lax about this.
- `Hello`
  - More than one `Hello` command was sent. Only one should be sent, at the very beginning.
- An unknown command was sent (one whose `commandType` was not 0 or 1).

### 2.3.2 Client

On the client side, invalid uses of the protocol MUST be handled simply by disconnecting. This happens in the following situations:

- `Announce`
  - The server sends an `Announce` before the client has sent a `SetStation`.
- `Welcome`
  - The server sends a `Welcome` before the client has sent a `Hello`.
  - The server sends more than one `Welcome` at any point (not necessarily directly following the first `Welcome`).
- `InvalidCommand`
  - The server sends an `InvalidCommand`. This may indicate that the client itself is incorrect, or the server may have sent it out of error. In either case, the client MUST disconnect.
- An unknown response was sent (one whose `replyType` was not 0, 1, or 2).

## 3 Implementation Requirements

You MUST implement this project in C to help you become familiar with the Berkeley sockets API. If you are unfamiliar or rusty with C, read through the documentation linked on the course web page or contact the TAs for help. We will offer full language support and help with debugging tools.<sup>3</sup>

---

<sup>3</sup>You are allowed (and encouraged) to implement future projects in a language other than C, but please contact us beforehand to seek approval for your chosen language. When you do so, please link us to a web page containing the documentation for the language's networking library. You must not use high-level socket wrappers unless you write them yourself; we will tell you which libraries are and are not acceptable. Thus far, we have approved requests for C++, Scheme, Go, and Java (but still contact us if you want to use one of these). The TAs will offer limited language support, and your partner for the project must approve of your choice of language.

### 3.1 Correctness

You will write three separate programs, each of which will interact with the user to varying degrees. It is your responsibility to sanitize all input. In particular, your programs **MUST NOT** do anything which is disallowed by this specification, even if the user asks for it. The choice of how you deal with this (for example, displaying an error message to the user) is yours, but an implementation which behaves incorrectly, even if only when given incorrect input by the user, will be considered incorrect.

### 3.2 Clients

You will write two separate clients.

#### 3.2.1 UDP Client

The UDP client handles song data. The executable **MUST** be called `snowcast_listener`. Its command line **MUST** be:

```
snowcast_listener <udpport>
```

The UDP client **MUST** print all data received on the specified UDP port to `stdout`<sup>4</sup>.

#### 3.2.2 TCP Client

The TCP client handles the control data. The executable **MUST** be called `snowcast_control`. Its command line **MUST** be:

```
snowcast_control <servername> <serverport> <udpport>
```

`<servername>` represents the IP address (*e.g.* 128.148.38.158) or hostname (*e.g.* localhost, cslab6c) which the control client should connect to, and `<serverport>` is the port to connect to. `<udpport>` is the port on which the local UDP client is watching for song data.

The control client **MUST** connect to the server and communicate with it according to the protocol. After the handshake, it **MUST** show a prompt and wait for input from `stdin`. If the user types in 'q' followed by a newline, the client **MUST** quit. If the user types in a number followed by a newline, the control **MUST** send a `SetStation` command with the user-provided station number, unless that station number is outside the range given by the server; you may choose how to handle this situation.

If the client gets an invalid reply from the server (one whose `replyType` is not 0, 1, or 2), then it **MUST** close the connection and exit.

The client **MUST** print whatever information the server sends it (*e.g.* the `numStations` in a `Welcome`). It **MUST** print replies in real time.

---

<sup>4</sup>There's no need for the UDP client to play the data it receives itself, since you can just pipe its output into another program which plays the music instead. More on this later.

### 3.3 Server

The server executable **MUST** be called `snowcast_server`. Its command line **MUST** be:

```
snowcast_server <tcpport> <file>...
```

That is, a port number on which the server will listen, followed by a list of one or more files. To make things easy, each station will contain just one song. Station 0 plays the first file, Station 1 plays the second file, etc... Each station **MUST** loop its song indefinitely.

When the server starts, it **MUST** begin listening for connections. When a client connects, it **MUST** interact with it as specified by the Protocol. Additionally, it **MUST** send an **Announce** whenever a song repeats.

You want the server to stream music, not to send it as fast as possible. Assume that all mp3 files are 128 Kibps, meaning that the server **MUST** send data at a rate of 128Kibps (16 KiB/s).<sup>5</sup>

The server **MUST** print out any commands it receives and any replies it sends to stdout. It will also have a simple command-line interface: ‘p’ followed by a newline **MUST** cause the the server to print out a list of its stations along with the clients that are connected to each one, and ‘q’ followed by a newline **MUST** cause the server to close all connections, free any resources it’s using, and quit.

Additionally:

- The server **MUST** support multiple clients simultaneously.
- There **MUST** be no hard-coded limit to the number of stations your server can support or to the number of clients connected to a station.
- Remember to properly handle invalid commands (see the Protocol section above).
- The server **MUST** never crash, even when a misbehaving client connects to it. The connection to *that* client **MAY** be terminated, however.
- If multiple clients are connected to one station, they **MUST** all be listening to the same part of the song, even if they connected at different times.
- If no clients are connected to a station, the current position in the song **MUST** still progress, without sending any data. The radio doesn’t stop when no one is listening.
- The server **MUST NOT** simply read the entire song file into memory at once. It **MAY** read the entire file in for some sizes, but there must be a size beyond which it will be read in chunks.<sup>6</sup>

## 4 Testing

We’ve provided a sample Makefile in `/course/cs168/pub/snowcast/Makefile` that you can use as a stencil to get started.

---

<sup>5</sup>You can convert an existing MP3 with: `ffmpeg -b 128k -i old.mp3 new.mp3`

<sup>6</sup>For example, you should be able to have `/dev/urandom` as a station

A good way to test your code at the beginning is to stream text files instead of mp3s. Once you're more confident of your code, you can test your program using the mp3 files in `/course/cs168/pub/snowcast/mp3`. You can pipe the output of your UDP client into `mpg123` to listen to the mp3:

```
./snowcast_listener port | mpg123 -
```

If you bring headphones to the Sunlab, you should hear something.

## 4.1 Rate Monitor

Unfortunately, there are many details to streaming mp3s well that would require understanding the mp3 file format in detail to do a really good job. Instead we ask only that you stream the mp3 at a constant bit rate. We've created a rate monitoring program available in `/course/cs168/pub/tools/<os>/rate_monitor`, where `<os>` is one of `linux`, `darwin`, or `windows`. This takes data from `stdin`, outputs it to `stdout`, and prints statistics about the rate at which it is receiving data to `stderr`. We'll be testing to see that your rate is consistently 16 KiB/second. You can run it as follows:

```
./snowcast_listener port | /course/cs168/pub/tools/<os>/rate_monitor > /dev/null
```

You can also run it with the `-h` flag to see more options. You can also pipe the rate monitor's output into `mpg123`.

## 4.2 Reference Implementations

For your convenience, we have provided binaries of reference implementations of the client and the server that follow the protocol and meet all the requirements. They're in `/course/cs168/pub/<os>/snowcast`, where `<os>` is one of `linux`, `darwin`, or `windows`. Take advantage of these! You can test your adherence to the protocol based on how well your programs interact with them. This is why our protocol is specified so precisely. Your programs are expected to interoperate with ours.

# 5 Grading

## 5.1 Milestone - 30%

To make sure you're on the right track, 30% of your grade will be a milestone. It is strongly recommended that you begin working on the milestone as soon as possible.

You will handin your milestone before Thursday the 17th at 11:59pm. In addition, you will have to schedule an appointment with a TA by Thursday the 17th at 6PM.

Within the milestone, 10% of the grade is a small demo. You must demo a fully functioning client that connects to our server, completes the handshake, and implements a read-evaluate-print-loop (REPL) that allows users to select a station, logs server communication, and quits. The client does NOT have to handle invalid server communications or other invalid conditions for the milestone.

The other 20% is for the design of your server, which will help you prepare for the hardest part of the assignment. You will be graded based on how well you have thought your design through. Your thread model is especially important. Consider these questions:

- Will you spawn a new thread for every command you receive?
- How many threads will you have per client?
- Will you have one thread to handle all the stations, or one or more threads for each one?
- How will threads communicate with each other? What mutexes will you need?
- How will you enforce that each client maintains proper use of the protocol?

If you're having trouble with the design, please come to our hours, or e-mail us with your question. We also encourage appointments outside of our hours if you feel you need help in-person.

## 5.2 Program - 65%

Most of your grade will be based on how well your program conforms to the specification. This includes how well it interacts with the reference implementations, as well as with each other's projects. Furthermore:

- You **MUST** check return codes for all system calls you make. You can use `perror` to print error messages.
- You can't assume `recv` and `send` will read or write all the bytes you requested. You have to check each return code and re-call them until the entire buffer is read or written.
- You **MUST** protect access to data shared by multiple threads, even integers.
- You **MUST** correctly handle clients or servers that do not adhere to the protocol by following the Invalid Commands section.

## 5.3 README - 5%

Please include a README file with your program. Describe design decisions, such as how your server is structured in terms of threading, how it handles announces, how it handles multiple clients, etc. List any bugs that you know your program has. We'll take off less points for any bugs you list than if we had to find them ourselves :).

# 6 Handin and Interactive Grading

For both the milestone and the final submission, hand in your project by typing

```
cs168_handin snowcast
```



from *inside* the directory where your work is located. To reduce clutter, the handin script removes .o files and binary executable files, and runs `make clean` before handing in your assignment. You can handin more than once - the new handin will replace the older one. We should be able to rebuild your programs by running `make`.

The milestone must be handed in by Thursday the 17th and the final assignment must be handed in by Thursday the 24th.

You will schedule an appointment to meet with the TAs for your interactive grading session in order to demonstrate the functionality of your client during the milestone, and for both your client and server for the final submission. Between the time you've handed in and each demo meeting, you can continue to make minor tweaks and bug fixes. However, the version you've handed in should be nearly complete since it could be referenced for portions of the grading.

## A Useful Hints/Tips

For the TCP connection, use `recv()` and `send()` (or `read()` and `write()`). For the UDP connection, use `sendto()` and `recvfrom()`. Don't send more than 1400 bytes with one call to `sendto()`<sup>7</sup>

For the TCP connection, timeouts can be set on the socket using `setsockopt()`.

You will want to permit the server to reuse its port, so that you can kill it and restart it without waiting a few minutes. Look at the end of section 4.2 in the networking guide (off the course website). To handle multiple connections on the server, you should have a thread which calls `accept()` in a loop. When it accepts a connection, it should start one or more threads to handle that connection, and then continue `accept()`ing. To control the rate that the server sends song data at, use the `nanosleep()` and `gettimeofday()` functions.

The TCP client has to read input from two sources at the same time - `stdin`, and the server. You might do this with a thread for the server and a thread for standard input, or you might use `select()`<sup>8</sup> to handle both tasks in a single thread without blocking. To implement hostname lookup (*e.g.* `localhost` to `127.0.0.1` or `cslab6e` to `128.148.31.38`), use `gethostbyname()`.

---

Please let us know if you find any mistakes, inconsistencies, or confusing language in this or any other CS168 document by filling out the anonymous feedback form:

<http://cs.brown.edu/courses/cs168/f15/feedback.html>.

---

<sup>7</sup>This is because the MTU of Ethernet is 1440 bytes, and we don't want our UDP packets to be fragmented. You'll learn more about this later.

<sup>8</sup>See <http://www.lowtek.com/sockets/select.html> for a guide on `select()`