

15-441 Computer Networks

Project 1: IRC and Routing

Assigned: Thursday, January 20, 2005

Due: Thursday, February 24, 2005 (11:59 PM EST)

Lead TAs: Maksim Tsvetovat <maksim@cs.cmu.edu>
David Craft <dcraft+441@andrew.cmu.edu>

1 Introduction

The purpose of this project is to give you experience in developing concurrent network applications. You will use the BERKELEY SOCKETS API to write an Internet chat server using a subset of the Internet Relay Chat protocol (IRC)¹ and implement two different routing protocols so chat messages can be exchanged between a network of chat servers.

IRC is a global, distributed, real-time chat system that operates over the Internet. An IRC network consists of a set of interconnected servers. Once users are connected to an IRC server, they can converse with other users connected to any server in the IRC network. IRC provides for group communication, via named channels, as well as personal communication through “private” messages. For more information about IRC, including available client software and public IRC networks, please see The IRC Prelude.²

If you have not used IRC before, you may want to try it out to get a feel for what it is. For a quick start, log in to an Andrew machine, and run `irc <nickname> irc.debian.org` where <nickname> is the nickname you want to use. Then type `/join #c<cr>` to join the C programming channel. Other channels you might be interested include `#debian`, `#redhat`, `#perl`, `#Linux`, and `#c++`. After you have tried out the text mode IRC client, you may want to try out more elaborate clients, such as `xchat` and `chatzilla` (part of mozilla).

In addition, you will also implement a shortest path link state routing protocol. In this protocol, each node in the network periodically exchanges information with its neighbors so that everyone in the network knows the best path to take to reach each destination. This is similar to the protocols used in Internet routers. At the end of this project, you will have your own network of chat servers, which could be used to talk with users across the world.

2 Logistics

- The tar file for this project can be found at:

<http://www.cs.cmu.edu/~srini/15-441/S05/assignments/project1/project1.tar.gz>

- This is a group project. You **must** find exactly one partner for this assignment. The only reason you should not have a partner is if there are an odd number of people in the class and you are left out (in which case contact us). Talk to your neighbors and use the bboards.
- Once you have found a partner, email David (dcraft+441@andrew.cmu.edu) your names and andrew logins so we can assign a group number to you. Use “**15441 GROUP**” as the subject line. Please try to be sure you know who you will work with for the full duration of the project so we can avoid the hassle of people switching later.
- This is a large project, but not impossible. Here is a recommended set of suggested milestones:

¹ <http://www.irchelp.org/irchelp/rfc/>

² <http://www.irchelp.org/irchelp/new2irc.html>

√	Date	Milestone
	1/20	Project Assigned
	1/23	Read and understand project handout and RFC
	1/31	Simple standalone IRC server complete
	2/1	CHECKPOINT open for business! (see boards)
	2/4	Simple standalone IRC server tested thoroughly
	2/7	CHECKPOINT DUE (may do earlier!)
	2/14	Routing daemon implementation complete
	2/18	Routing daemon tested thoroughly
	2/20	IRC server extensions / forwarding protocol complete
	2/23	IRC server extensions tested thoroughly
	2/24	Last minute rush to get things done and hand-in

3 General Overview

An IRC network is composed of a set of nodes interconnected by virtual links in an arbitrary topology. Each node runs a process that we will call a routing daemon. Each routing daemon maintains a list of IRC users available to the system. Figure 1 shows a sample IRC network composed of 5 nodes. The solid lines represent virtual links between the nodes. Each node publishes a set of users (i.e., the nicks of the IRC clients connected to it) to the system. The dotted lines connect the nodes to their user sets.

The usage model is the following: If Bob wants to contact Alice, the IRC server on the left first must find the route or path from it to the node on the right. Then, it must forward Bob's message to each node along the path (the dashed line in the figure) until it reaches the IRC server at Alice's node, which can then send the message to the client Alice.

In essence, each node in the system performs functions similar to the ones performed in the network layer, namely forwarding and routing. Forwarding is the action performed by each node to guide a packet toward its destination. Routing refers to the action of building the data structures necessary to reach particular destinations (in terms of the IRC server, a destination is a username/nick).

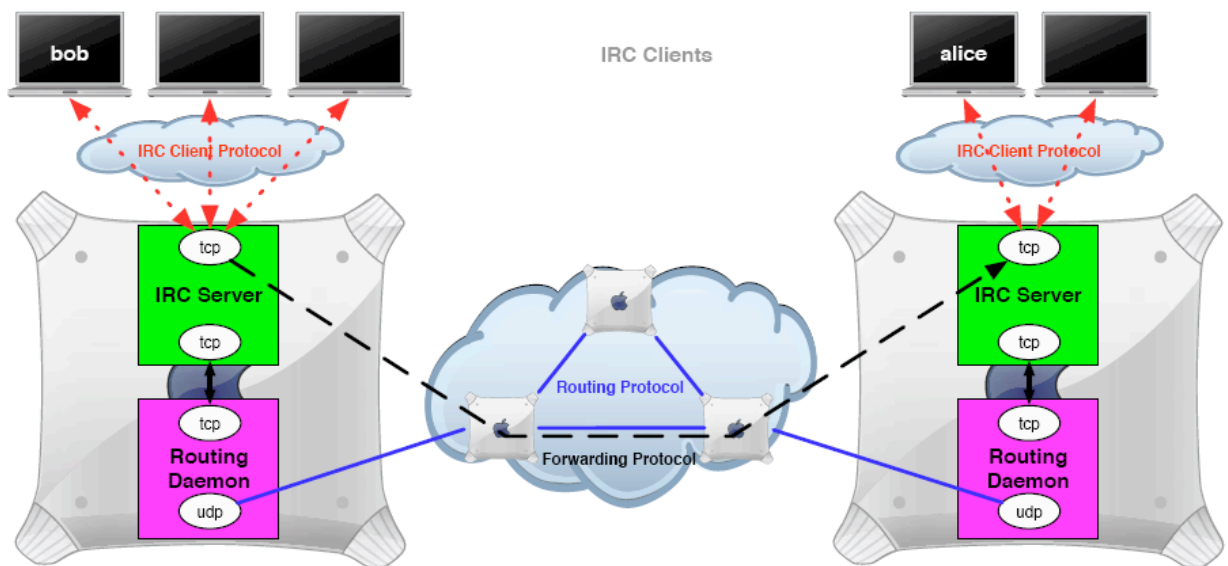


Figure 1 – Sample IRC Network

The routing daemon will be a separate program from your IRC server. Its purpose is to maintain the routing state of the network (e.g., build the routing tables or discover the routes to destinations). When the IRC server wants to send a message to a remote user, it will ask the routing daemon how to get there and then send the message itself. In other words, the routing daemon does the routing and the IRC server does the forwarding.³

In your implementation, the routing daemon will communicate with other routing daemons (on other nodes) over a UDP socket to exchange routing state. It will talk to the IRC server that is on the same node as it via a local TCP socket. The IRC server will talk to other IRC servers via the TCP socket that it also uses to communicate with clients. It will simply use special server commands. This high level design is shown in the two large IRC server nodes in Figure 1.

In order to find out about the network topology, each routing daemon will receive a list of neighboring nodes when it starts. In this project, you can assume that the no new nodes or links will ever be added to the topology after starting, but nodes and links can fail (i.e., crash or go down) during operation (and may recover after failing).

4 Definitions

Before jumping into the gory details, let us define some terminology.

- **node** – An IRC server and routing daemon pair running together that is part of the larger network. In the real world, a node would refer to a single computer, but we can run multiple “virtual” nodes on the same computer since they can each run on different ports. Each node is identified by its `nodeID`.
- **nodeID** – A unique identifier that identifies a node. This is an unsigned 32-bit integer that is assigned to each node when its IRC server and routing daemon start up.
- **neighbor** – Node 1 is a neighbor of node 2 if there is a virtual link between 1 and 2. Each node obtains a list of its neighbors’ `nodeIDs` and their routing and forwarding ports at startup.
- **destination** – An IRC username or nick as a null terminated character string. As per the IRC RFC, destinations will be at most 9 characters long and may not contain spaces.
- **IRC port** – The TCP port on the IRC server that talks to clients and other IRC servers.
- **forwarding port** – Same as IRC port.
- **routing port** – The UDP port on the routing daemon used to exchange routing information with other routing daemons.
- **local port** – The TCP port on the routing daemon that is used to exchange information between it and the local IRC server. For example, when the IRC server wants to find out the route to remote user, it queries the routing daemon on this port. The socket open for listening will be on the routing daemon. The IRC server will connect to it.
- **OSPF** – The shortest path link state algorithm you will implement
- **routing table** – The data structure used to store the “next hops” that packet should take used in OSPF. See your textbook pp. 274 – 280 for description.

5 The IRC Server

³ In actual routers, and even overlay networks like peer-to-peer file sharing networks, the notion of the separate routing daemon is atypical. Normally, the forwarding program should keep the forwarding table, not query the route daemon for each route lookup.

Your server will implement a subset of the original IRC protocol. The original IRC protocol is defined in RFC 1459⁴. Because RFC 1459 omits some details that are required to implement an IRC server, we have provided an annotated version of the RFC⁵. For this project, **you should always refer to the annotated version of the RFC**, not the original version. We have chosen a subset of the protocol that will provide you with experience developing a concurrent network application without spending an inordinate amount of time implementing lots of features. Specifically, your server must implement the following commands:

Basic Commands:

- **NICK:** Give the user a nickname or change the previous one. Your server should handle duplicate nicknames and report an error message.
- **USER:** Specify the username, hostname, and real name of a user.
- **QUIT:** End the client session. The server should announce the client's departure to all other users sharing the channel with the departing client.

Channel Commands:

- **JOIN:** Start listening to a specific channel. Although the standard IRC protocol allows a client to join multiple channels simultaneously, your server should restrict a client to be a member of at most one channel. Joining a new channel should implicitly cause the client to leave the current channel.
- **PART:** Depart a specific channel. Though a user may only be in one channel at a time, **PART** should still handle multiple arguments. If no such channel exists or it exists but the user is not currently in that channel, send the appropriate error message.
- **LIST:** List all existing channels on the local server only. Your server should ignore parameters and list all channels and the number of users on the local server in each channel.

Advanced Commands:

- **PRIVMSG:** Send messages to users. The target can be either a nickname or a channel. If the target is a channel, the message will be broadcast to every user on the specified channel, except the message originator. If the target is a nickname, the message will be sent only to that user.
- **WHO:** Query information about clients or channels. In this project, your server only needs to support querying channels on the local server. It should do an exact match on the channel name and return the users on that channel.

For all other commands, your server must return `ERR_UNKNOWNCOMMAND`. If you are unable to implement one of the above commands (perhaps you ran out of time), your server must return the error code `ERR_UNKNOWNCOMMAND`, rather than failing silently, or in some other manner.

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` — typically 1024). 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 up if a client sends only a partial command. 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.** 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 multiple commands in one packet. It should not overflow any buffers when the

⁴ <http://www.ietf.org/rfc/rfc1459.txt>

⁵ <http://www.cs.cmu.edu/~srini/15-441/S05/assignments/project1/rfc-html/index.html>

client sends a message that is too long (longer than 512 bytes). In general, your server should not be vulnerable to a malicious client.

Note your server behaves differently from a standard IRC server for some of the required commands (e.g., JOIN). Therefore, you should not use a standard IRC server as your reference for your implementation. Instead, refer to the annotated version of the RFC on the course web page.

6 Link-State Routing Protocol

6.1 Basic Operation

You will implement a link-state routing protocol similar to OSPF, which is described in the textbook in chapter 4, and in more detail in the OSPF RFC⁶. Note, however, that your protocol is greatly simplified compared to the actual OSPF spec. As described in the references, OSPF works by having each router maintain an identical database describing the network's topology. From this database, a routing table is calculated by constructing a shortest-path tree. Each routing update contains the node's list of neighbors, users, and channel. Upon receiving a routing update, a node updates its routing table with the "best" routes to each destination. In addition, each routing daemon must remove entries from its routing table when they have not been updated for a long time. The routing daemon will have a loop that looks similar the following:

```
while (1)
{
    /* each iteration of this loop is "cycle" */
    wait_for_event(event);

    if (event == INCOMING_ADVERTISEMENT)
    {
        process_incoming_advertisements_from_neighbor();
    }
    else if (event == IT_IS_TIME_TO_ADVERTISE_ROUTES)
    {
        advertise_all_routes_to_all_neighbors();
        check_for_down_neighbors();
        expire_old_routes();
    }
}
```

Let's walk through each step. First, our routing daemon A waits for an event. If the event is an incoming link-state advertisement (LSA), it receives the advertisement and updates its routing table if the LSA is new or has a higher sequence number than the previous entries. If the routing advertisement is from a new router B or has a higher sequence number than the previously observed advertisement from router B, our router A will flood the new announcement to all of its neighbors except the one from which the announcement was received, and will then update its own routing tables.

If the event indicates that a predefined period of time has elapsed and it is time to advertise the routes, then the router advertises all of its users, channels, and links to its direct neighbors. If the routing daemon has not received any such advertisements from a particular neighbor for a number of advertisements, the routing daemon should consider that neighbor down. The daemon should mark the neighbor down and re-flood LSA announcements from that neighbor with a TTL of zero. When your router receives an announcement with a TTL of zero, it should delete the corresponding LSAs from its

⁶ <http://www.rfc-editor.org/rfc/rfc2328.txt>

table.

If the event indicates that a user has joined or left a channel or the server, the router should send a *triggered update* to its neighbors. This is simply a new link state advertisement with a higher sequence number that announces the router's new state. If a node has not sent any announcements for a very long time, we expire it by removing it from our table.

If B receives an LSA announcement from A with a lower sequence number than it has previously seen (which can happen, for example, if A reboots), B should echo the prior LSA back to A. When A receives its own announcement back with a higher sequence number, it will increment its transmitted serial number to exceed that of the older LSAs.

Each routing announcement should contain a full state announcement from the router – all of its neighbors, all of its users, and all of its channels. This is an inefficient way to manage the announcements (see the extra credit section), but it greatly simplifies the design and implementation of the routing protocol to make it more tractable for a 5 week assignment. Each time your router originates a new LSA, it should increment the serial number it uses. When a router receives an updated LSA, it recomputes its local routing table. The LSAs received from each of the peer nodes tell the router a link in the complete router graph. When a router has received all of the LSAs for the network, it knows the complete graph. Generating the user routing table is simply a matter of running a shortest-paths algorithm over this graph.

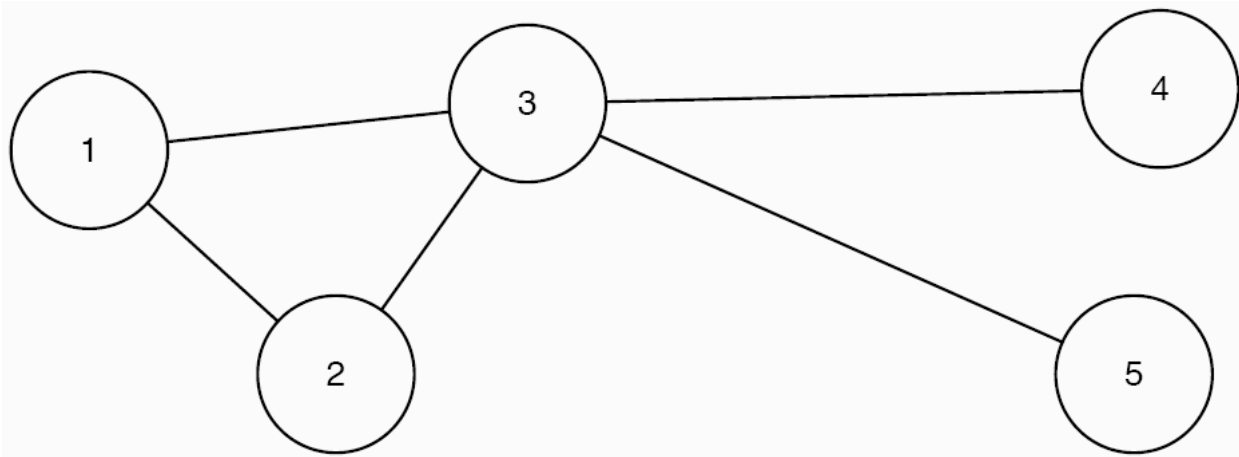


Figure 2 – Sample Network

6.2 Reliable Flooding

OSPF is based upon reliable flooding of link-state advertisements to ensure that every node has an identical copy of the routing state database. After the flooding process, every node should know the exact network topology. When a new LSA arrives at a router, it checks to see if the sequence number on the LSA is higher than it has seen before. If so, the router *reliably* transmits the message to each of its peers except the one from which the message arrived.

The flooding is made reliable by the use of acknowledgement packets from the neighbors. When router A floods an LSA to router B, router B responds with an “LSA Ack.” If router A does not receive such an ack from its neighbor within a certain amount of time, router A will retransmit the LSA to B.

With the information contained in the LSAs, each server should be able to deliver messages from one user to another without much trouble. To send messages to a channel, however, requires a little more work; this is multicast routing instead of unicast routing. A channel can exist on multiple servers, so the distribution can take multiple branches at a time. How does the local node know which neighbors to forward the message to in this case?

Since a channel can exist on multiple servers and the server knows the network topology, each server is able to construct a source rooted shortest paths tree for that message, which tells the server what

outgoing links it should use. Note that this tree is rooted at the *message* source, not the router making the computation. The algorithm for computing the shortest paths tree for the multicast case is the same as for the unicast to a user case, except that the source may not be the local node.

Why does it need to know the source? Consider the network pictured in Figure 2. Now, suppose nodes 1, 3, and 4 have users subscribed to channel #perl, and nodes 2, 3, and 5 have users subscribed to channel #c. If a user on node 1 wants to send a message to #perl then it should propagate the message to node 3. Node 3 knows nodes 1 and 4 also have users in #perl, but since the message came from node 1, it should not propagate the message back to node 1. So, it only propagates the message to node 4.

The channel routing table for node 3 would look like:

Channel name	Source Node	Next-Hops
#perl	1	4
#perl	2	4
#perl	3	1, 4
#perl	4	1
#perl	5	1, 4
#c	1	5
#c	2	5
#c	3	2, 5
#c	4	2, 5
#c	5	2

The channel routing table for node 1 would look like:

Channel name	Source Node	Next-Hops
#perl	1	3
#perl	2	None
#perl	3	None
#c	1	2, 3
#c	2	None
#c	3	None

The channel routing table for node 5 would look like:

Channel name	Source Node	Next-Hops
#perl	3	None
#perl	5	3
#c	3	None
#c	5	3

There are two ways that a router could potentially compute these routes. It could compute all possible (source, channel) trees in advance, and populate its routing table with the results. Or, the router could compute the routes dynamically on demand when a message arrives for a particular channel, and cache the results. **In this assignment, we will implement the dynamic computation and caching version.** The router's multicast routing table (channel routing table) should act as a cache, and the router should compute the trees dynamically if it has no entry. Remember to have a way to indicate "route calculated, but there were no neighbors" in your routing table so that you don't eternally recompute local channels.

Your router should be robust to misbehaving neighbors. If the router receives a channel message forwarded to it by a peer who should not have forwarded the message (remember, the router can check

this, since it knows the shortest paths tree), it should silently drop the message. Such an event could occur during a route change when the routing table became temporarily inconsistent, and it can lead to routing loops. Because multicast can generate a large number of packets, and the IRC network is a less trusted environment than an ISP's own routers, dropping the message is safer (for the project), but might cause packet delivery to be somewhat less reliable. The congestion caused by routing loops is also typically addressed through a TTL in the packets being forwarded.

6.3 Protocol Specifications

Figure 3 shows the routing update message format, with the size of each field in bytes in parenthesis.

Version (1), TTL (1), type (2)
sender nodeID (4)
sequence number (4)
num link entries (4)
num user entries (4)
num channel entries (4)
Link entries (variable)
User entries (variable)
Channel entries (variable)

Figure 3 – Link State Packet Format

- **Version** – the protocol version, always set to 1
- **TTL** – the time to live of the LSA. It is decremented each hop during flooding, and is initially set to 32.
- **Type** –Announcement packets should be type 0. Acknowledgements should be type 1.
- **Sender nodeID** – The nodeID of the sender, so the receiver knows which neighbor it came from.
- **Num link entries** – The number of link table entries in the message.
- **Sequence number** – The sequence number for the LSA.
- **Num user entries** – The numbers of users announced in the message
- **Num channel entries** – The number of channels announced in the message
- **Link entries** – Each link entry contains the nodeID of a node that is directly connected to the sender. This field is 4 bytes.
- **User entries** – Each user entry contains the name of the destination user as a null terminated string. Since the IRC RFC indicates that nicknames should be at most 9 characters and we have added the constraint that channels can be at most 9 characters (including & or #), it should definitely fit within 16 (the unused bytes will be ignored).
- **Channel entries** – Same as a user entry, above.

All multi-byte integer fields (nodeIDs, TTLs, link entries, the type field, etc) should be in **network byte order**.

An acknowledgement packet looks very similar to an announcement packet, but it does not contain any entries. It contains the sender nodeID and sequence number of the original announcement, so that the peer knows that the LSA has been reliably received.

6.4 Requirements

Your implementation of the link-state routing protocol should have the following features:

- Given a particular network configuration, the routing tables at all nodes should converge so that forwarding will take place on the paths with shortest length.
- In the event of a tie for shortest path, the next hop in the routing table should always point to the nodeID with the lowest numerical value. Note that this implies there should be a unique solution to the routing tables in any given network.
- Remove the LSAs for a neighbor if it hasn't given any updates for some period of time.
- You should implement *Triggered Updates* (when a link goes down or when users join or leave a server or channel).
- If a node or link goes down (e.g., routing daemon crashes, or link between them no longer works and drops all messages), your routing tables in the network should re-converge to reflect the new network graph. You shouldn't have to do anything more to make sure this happens, as the above protocol already ensures it.

You do not have to implement the following:

- You do not have to provide authentication or security for your routing protocol messages.
- You only need to store the single best route to a given user.
- You do not have to “jitter” your timer with randomized times.

The choice of shortest-paths algorithm is up to you; we suggest Dijkstra's. You don't need to implement a highly efficient priority queue for this assignment; use whatever tools (STL, etc.) that you're familiar with.

7 Local Server–Daemon Protocol

This section describes the mini-protocol that an IRC Server uses to talk to the local routing daemon on the same node. It is important that you follow these specifications carefully because we will test your routing daemon independently of your IRC server!

The routing daemon listens on the local port when it starts up to service route lookup requests. When the IRC server on the same node starts up, it connects to the local port of the routing daemon. Since the local port is only supposed to service local client programs (like the IRC server) on the same machine that it trusts, you can assume that we won't do anything intentionally malicious to try to break it. However, you may find it useful to make it robust to invalid input, since you may make typos when testing it. Specifically, you can assume:

- We will only use the protocol as defined below. We will not send invalid requests.
- Only a single IRC server will connect to the routing daemon.
- Your IRC server may block while waiting for a response from the routing daemon. (i.e., you can treat it as a function call)

This is a line-based protocol like the IRC-protocol itself.
Each request and response pair looks like this:

command arguments . . .
results . . .

Where *command* is the name of the request, *arguments . . .* is a space-separated list of arguments to the command, and *results . . .* is a space-separated list of results returned. All requests and responses are terminated with a newline character (\n) and are case sensitive, but some responses have multiple lines. You must implement the following request/response pairs in your routing daemon:

Request: ADDUSER *nick*

Response: OK

Description: This request is issued when a new user is registered with the IRC server. The user's nick is added to the routing daemon's list of local users so that other nodes can find the user. This should trigger an immediate update for that nick.

Examples:

req: ADDUSER bob
resp: OK
req: ADDUSER alice
resp: OK

Request: ADDCHAN *channel*

Response: OK

Description: This request is issued when a new user is registered with the IRC server. The channel name is added to the routing daemon's list of local channels so that other nodes can find the channel. This should trigger an immediate update for that channel.

Examples:

req: ADDCHAN #perl
resp: OK
req: ADDCHAN #networks
resp: OK

Request: REMOVEUSER *nick*

Response: OK

Description: This request is issued when a local user leaves the IRC server. The user's nick is removed from the routing daemon's list of local destinations so that other nodes will know that they can no longer reach the user there. This should trigger an immediate update for that nick.

Examples:

req: REMOVEUSER bob
resp: OK
req: REMOVEUSER baduser
resp: OK

Request: REMOVECHAN *channel*

Response: OK

Description: This request is issued when the last local user leaves a channel. The channel name is removed from the routing daemon's list of local channels so that other nodes will know that they should no longer send channel messages to that server. This should trigger an immediate update for that channel.

Examples:

```
req:    REMOVECHAN bob
resp:   OK
req:    REMOVECHAN baduser
resp:   OK
```

Request: NEXTHOP *nick*

Response: OK *nodeID distance*

Response: NONE

Description: This request is used to find nodeID of the next hop to use if we want to forward a message to the user *nick*. It should return OK if the routing table has a valid next hop for the *nick* along with the distance to that destination, and NONE otherwise (e.g., if the destination's distance is not known or user does not exist).

Examples:

```
req:    NEXTHOP bob
resp:   OK 2 5
req:    NEXTHOP alice
resp:   OK 3 2
req:    NEXTHOP baduser
resp:   NONE
```

Request: NEXTHOPS *sourceID channel*

Response: OK *nodeID nodeID nodeID ...*

Response: NONE

Description: This request is used to find which links a server should send messages to if it wants to forward a message to a channel. It should return OK if the routing table has a valid entry for the channel from the given source node and then list the nodes to which it should propagate the message. Otherwise, it should return NONE (e.g., if the channel does not exist, or only exists on the local node). See graph in Link-State section.

Examples:

```
req:    NEXTHOPS 1 #perl
resp:   OK 2 5 9
req:    NEXTHOPS 5 #perl
resp:   OK 1
req:    NEXTHOPS #badchan
resp:   NONE
```

Request: USERTABLE

Response: OK *size*

Description: If this request is issued, the routing daemon should respond with OK, the *size* or number of entries in the routing table, and a multi-line response with its entire user table in the following format:

```
nickname next-hop distance
nickname next-hop distance
nickname next-hop distance
...
```

Where *nick* is the nickname, *next-hop* is the nodeID of the next hop, and *distance* is the current distance value for that destination. You should not include local nicknames in this list. The order of entries does not matter. Your IRC Server will probably not need to use this command. We will use this to test your routing daemon. This would be similar to calling NEXTHOP on every user on the server.

Examples:

```
req:    USERTABLE
resp:   OK 3
        bob 2 2
        alice 3 1
        jim 3 2
```

Request: CHANTABLE

Response: OK *size*

Description: If this request is issued, the routing daemon should respond with OK, the *size* or number of entries in the channel table, and a multi-line response with its entire channel table in the following format:

```
channel sourceID next-hop next-hop next-hop ...
channel sourceID next-hop next-hop next-hop ...
channel sourceID next-hop next-hop next-hop ...
...
```

Where *channel* is the channel name, *sourceID* is the nodeID on which the message would come, and *next-hop* is a list of nodeIDs to which the server should propagate a message for that channel. You should not include channels that exist only locally in this list. The order of entries does not matter. Your IRC Server will probably not need to use this command. We will use this to test your routing daemon. This would be similar to calling NEXTHOPS on every channel on the server.

Examples:

```
req:    CHANTABLE
resp:   OK 3
        #perl 1 2 5 9
        #perl 2 5 9
        #perl 5 1 9
        #perl 9 5 2
```

8 IRC Server (revisited)

Now that we have covered the IRC server, the routing protocols, and the server-daemon protocol, the only major issue remaining is how to extend your IRC Server to use the routing daemon so it can send messages to users on remote IRC Servers.

Remember that the PRIVMSG command has two targets: nicknames and channels. If the target is a nickname, the IRC server must first determine if there is a local user with that nickname. If not, then it should try to locate the user on a remote IRC Server (using the routing daemon) and, if found, forward the

message to that IRC Server which will then send it to the target. If the target is not found, then you should send the user an `ERR_NOSUCHNICK` error. If the target is a channel, then you must echo that message to every user on that channel.

8.1 Requirements

Your extensions to the IRC server should have the following features:

- Connect to the routing daemon's local port when it starts up. You can assume the routing daemon will be started first.
- When a new user is registered with the IRC server, it should add the user's nick to the routing daemon's list of users using the `ADDUSER` request.
- When a user leaves the IRC server, it should remove the user's nick from the routing daemon's list of users using the `REMOVEUSER` request.
- When a channel is created on the IRC server, it should send an `ADDCHAN` message to the routing daemon.
- When the last user leaves a channel on the IRC server, it should send a `REMOVECHAN` message to the routing daemon.
- If a user changes his or her nick, remove the old nick and add the new one to the routing daemon.
- When a `PRIVMSG` is sent to a nick that we don't know locally, the IRC Server should ask the routing daemon to find it, if possible, and forward the message to that user. The remote IRC server receiving the message should send it to the target user the same way it would send any other `PRIVMSG` to him or her.
- If the target is not found, then you should send the user an `ERR_NOSUCHNICK` error as defined in section 4.4.1 of the IRC RFC.
- The `PRIVMSG` command should support multiple targets; i.e., the `PRIVMSG` command may have a comma-separated list of target users or channels that should all be sent the message.
- If the routing daemon dies or you cannot communicate with it, your IRC server may exit.

You do not have to implement the following:

- Forwarding messages to target servers, host masks, or anything mentioned in the IRC RFC that is not mentioned in this document.
- Forwarding any other type of message besides `PRIVMSG`.

8.2 Message Forwarding

Once the IRC Server has found the next hop or route to a remote nickname, it must forward the message to the remote IRC Server. You are responsible for designing a protocol to be used between your IRC Servers for forwarding these messages so that they will reach the destination. Here are a couple things to keep in mind when designing your protocol:

- When using OSPF, you can only obtain the next hop from the routing daemon. Hence, each IRC server along the path will have to query its routing daemon to figure out where to send the packet next.
- When using OSPF, while forwarding, a node or virtual link may go down (or the target user may leave). In this circumstance, you can just drop the message. You do not have to inform the user that sent the message that it was dropped.
- You may have to send the message to multiple peers when forwarding to a remote channel.
- If the same nick is logged on to more than one IRC Server in the network, OSPF should find the route to the "closest" one. Your forwarding protocol only needs to forward the message to one of

them.

- IRC Servers and virtual links may go down and come back up. If you detect that your neighbor is down (i.e., the socket is closed), you should check to see if they have come back up at least once every 3 seconds. In fact, when the network first starts up, since only one server will come up at a time, all its neighbors will appear to be down at first.
- You should not have IRC Servers communicate if they are not neighbors.
- Your forwarding protocol should not be “flood every message to every IRC server on the network.” That is not efficient and doesn’t require the routing layer at all.
- You should not rely on any special extensions to the local port mini-protocol. We may test your IRC Server on our own routing daemon.

9 Implementation Details & Usage

Your programs must be written in the C programming language. You are not allowed to use any custom socket classes or libraries, only the standard `libsocket`, the provided library functions, and the `csapp` wrapper library developed for 15-213. You may use the `pthread` library and STL, but you are responsible for learning how to use them correctly yourself if you choose to. To use the `csapp` wrapper library, you must link with `libpthread` (`-lpthread`). If you wish to use other libraries, please contact us.

9.1 Compiling

You responsible for making sure your code compiles and runs correctly on the Andrew x86 machines running Linux (i.e., `linux.andrew.cmu.edu`). We recommend using `gcc` to compile your program and `gdb` to debug it. You should use the `-Wall` flag when compiling to generate full warnings and to help debug. For this project, you will also be responsible for turning in a GNU Make (`gmake`) compatible Makefile. See the GNU make manual⁷ for details. When we run `gmake` we should end up with the routing daemon which you should call `srouded` and the simplified IRC Server which is called `sircd`.

9.2 Command Line Arguments

Your routing daemon must take the following command line arguments in any order. We will provide you some framework code that will read in these arguments.

usage: `./srouded -i nodeID -c config_file [options]`

-i *integer*

NodeID. Sets the nodeID for this process.

-c *filename*

Config file. Specifies the name of the configuration file that contains the information about the neighbor nodes. The format of this file is described below.

It should also recognize the following optional switches:

-a *seconds*

Advertisement cycle time. The length of time between each advertisement cycle. Defaults to 30.

⁷ http://www.gnu.org/manual/make/html_mono/make.html

-n *seconds*

Neighbor timeout. The elapsed time after which we declare a neighbor to be down if we have not received updates from it. You may assume that this value is a multiple of **advertisement cycle time**. Defaults to 120.

-r *seconds*

Retransmission timeout. The elapsed time after which a peer will attempt to retransmit an LSA to a neighbor if it has not yet received an acknowledgement for that LSA. This value is an integral number of seconds. Defaults to 3.

-t *seconds*

LSA timeout. The elapsed time after which we expire an LSA if we have not received updates for it. You may assume that this value is a multiple of **advertisement cycle time**. Defaults to 120.

Your IRC server will always have two arguments:

usage: `./sircd nodeID config_file`

nodeID

The nodeID of the node.

config_file

The configuration file name.

9.3 Configuration File Format

This file describes the *neighborhood* of a node. The neighborhood of a node 1 is composed by node 1 itself and all the nodes n that are directly connected to 1. For example, in Figure 4, the neighborhood of node 1 is {1, 2, 3}. The format of the configuration file very simple, and we will supply you with code to parse it. The file contains a series of entries, one entry per line. Each line has the following format:

```
nodeID hostname routing-port local-port IRC-port
```

nodeID

Assigns an identifier to each node.

hostname

The name or IP address of the machine where the neighbor node is running.

local-port

The TCP port on which the routing daemon should listen for the local IRC server.

routing-port

The UDP port where the neighbor node listens for routing messages.

IRC-port

The TCP port on which the IRC server listens for clients and other IRC servers.

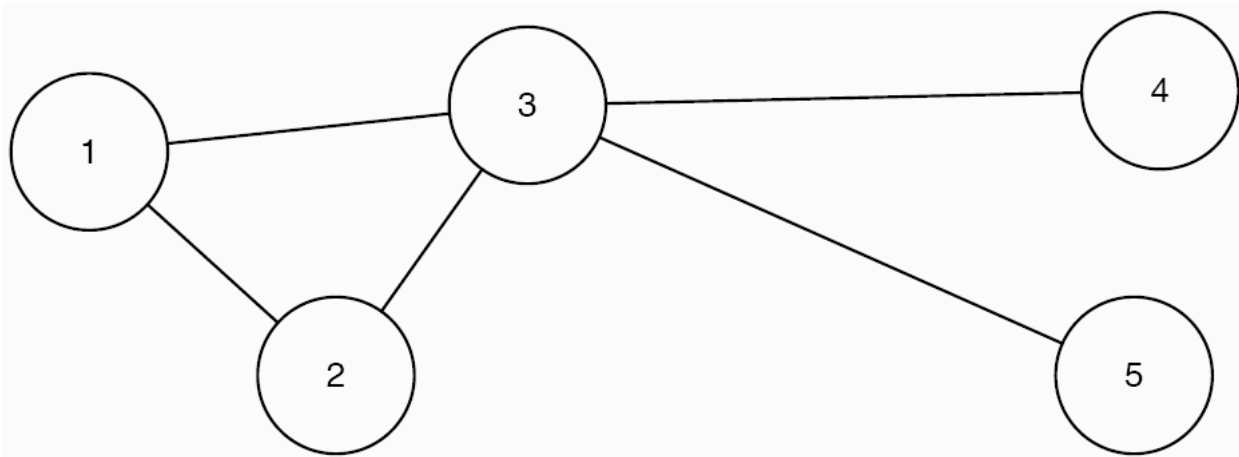


Figure 4 – Sample Node Network

Node 2	Node 5
<pre> 2 localhost 20203 20204 20205 1 unix1.andrew.cmu.edu 20200 20201 20202 3 unix3.andrew.cmu.edu 20206 20207 20208 </pre>	<pre> 3 unix3.andrew.cmu.edu 20206 20207 20208 5 localhost 20209 20210 20211 </pre>

Figure 5 – Sample Configuration file for nodes 2 and 5

How does a node find out which ports it should use as routing, IRC, and local ports? When reading the configuration file if an entry's `nodeID` matches the node's `nodeID` of the node (passed in on the command line), then the node uses the specified port numbers to route and forward packets. Figure 5 contains a sample configuration files corresponding to node 2 and node 5 for the network in Figure 4. Notice that the file for node 2 contains information about node 2 itself. Node 2 uses this information to configure itself.

We have provided you with a simple script called `genconfig.pl` that will auto-generate all the configuration files for a specified network graph, which you can find in the `./util` subdirectory of the handout. Read the text at the top of the script for documentation.

9.4 Running

This is how we will start your IRC network.

First, we start each routing daemon with the commands:

```

./sroute -i 0 -c node0.conf ... &
./sroute -i 1 -c node1.conf ... &
./sroute -i 2 -c node2.conf ... &
...

```

Each routing daemon will be started with its own configuration file to find out about its neighbors (described above) and its `nodeID`. In addition, we will pass it certain arguments to set the timer values.

Next, we will start each IRC server at each node:

```

./sircd 0 node0.conf &

```

```
./sircd 1 node1.conf &
./sircd 2 node2.conf &
...
```

Each IRC Server will be passed its nodeID and the configuration file to find out about its neighbors and what ports it should use/talk to. Now we will wait enough time such that the routing state should have converged and test your system. (We may also bring down nodes and restart them to test how resilient your system is to faults)

9.5 Framework Code

We have provided you with some framework code to simplify some tasks for you, like reading in the command line arguments and parsing the configuration file. You do not have to use any of this code if you do not want to. This code is documented in `rtlib.h` and implemented in `rtlib.c`. Feel free to modify this code also. However, you **must** use the following three routines, which are declared in `rtgrading.h` and implemented in `rtgrading.c`, and must **not** modify them:

- `rt_init(...)`: You must call this function when your routing daemon starts with the `argc` and `argv` passed to your program.
- `rt_sendto(...)`: Wrapper function for the `sendto()` system call. The parameters and semantics are the same as in the system call. You should use this function to send UDP packets in your routing daemon.
- `rt_recvfrom(...)`: Wrapper function for the `recvfrom()` system call. The parameters and semantics are the same as in the system call. You should use this function to receive UDP packets in your routing daemon.

We will replace `rtgrading.c` with implementations that we will use for grading so you should not modify it.

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 reading the b-boards 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 b-boards.

10 Testing

Code quality is of particular importance to 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 you know and acknowledge the pitfalls of your server, than miss them). Several paragraphs (or even a bulleted list of things done and why) should suffice for the test case documentation.

To help you get started on testing, we have provided a simple IRC client `sircc` and several

example test scripts. These will give you an idea of what tests we will use to evaluate your work, and ensure that you are on the right track with your server.

sircc:

The `sircc` program takes input from stdin as client commands to send to the server, and echoes server reply on the screen. This can be used to check the exact formats of responses from your server and test how your server behaves when given input is not compliant with the IRC specification.

```
unix>./sircc -h
usage: sircc <ip address> <port>
```

When using `sircc`, `<ip address>` and `<port>` are the address and port number of your IRC server. By default, the address is set to your local machine and the port number is 6667.

Ports:

To facilitate the development, you can run all your routing engines in the same machine using different ports. In order to avoid *port collisions* among routing engines from different groups, use port numbers in the range:

$$[20,000 + group\ number \times 100 , 20,000 + group\ number \times 100 + 99]$$

For example, if your group number is 23, then you should choose the port numbers from the following range: [22300, 22399]. If you are testing a configuration with 3 nodes, then the node addresses could be the following:

Node Id	IP Address	Routing Port	Local Port	IRC Port
1	127.0.0.1	22300	22301	22302
2	127.0.0.1	22303	22304	22305
3	127.0.0.1	22306	22307	22308

The utility scripts we provide do this port assignment for you. They can be found in `./util`

IRC Test scripts:

The test scripts test your IRC server against different types of commands.

For example, `login.exp` checks the replies of the command `NICK` and `USER`.

```
unix>./login.exp
usage: login.exp <host> <port>
```

Here `<host>` and `<port>` are the address and port number of your IRC server.

You may use the provided test scripts as a base to build your own test case. You may also find the following tools to be useful in your test code development:

irc, xchat, chatzilla

Various IRC clients. `irc` is a text mode client, while `xchat` and `chatzilla` have graphical interfaces. Since you are writing a genuine IRC server, you should be able to test it with actual IRC clients.

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.

Net::IRC

A Perl module that simplifies writing an IRC client. Net::IRC is not installed on the Andrew Linux machines, but you can download Net::IRC7 from the Comprehensive Perl Archive Network (CPAN).

Note that Net::IRC and a command line IRC client both implement the client-side IRC protocol for you. Presumably, they interact with the server in a standards-compliant manner.

Daemon Debugging:

The daemon will have no user interface, but you can still telnet to the local port on your routing daemons to inject destinations, remove destinations, check routing tables, etc.

To test if your system can handle node faults, kill some of your routing daemons and IRC servers. To test if your system can handle link faults, try blocking off a pair of UDP ports between two routing daemons (You can do this artificially in your code by dropping packets that go between them).

11 Handin

Your hand-in directory will be the following:

`/afs/andrew.cmu.edu/course/15/441/grp[your group #]`

You should submit the following files:

- **Makefile** – Make sure all the variables and paths are set correctly such that your program compiles in the handin directory. The Makefile should build two executable named `sroured` and `sircd`.
- 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 design of your routing daemon and a complete description of the protocols you used for forwarding IRC messages.
- **tests.txt**: File containing documentation of your test cases and any known issues you have.
- **extra.txt**: (optional) Documentation on any extra credit items you have worked on.

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

12 Grading

- **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 them (see the resources section below). Even if your server does not implement any IRC commands, your project submission can get up to 20 points here. Thus it is better to have partial functionality working solidly than lots of code that doesn’t actually do anything correctly.

- **Server IRC protocol:** 20 points

The grade in this section reflects how well you read, interpreted, and implemented the IRC protocol. We will test that all the commands specified in the project handout work. All commands 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. To receive full credit here, your server must accept a connection from `ircII`⁸ and be able to complete all required IRC commands as sent by this client.

- **OSPF User routing:** 25 points

The OSPF routing protocol should find a route if it exists. If there is more than one, it should only accept one and ignore the others. If there is no route, it should timeout after a specified time and ignore any path it might receive after timeout. If there are two users, you should use only one path and ignore the others.

- **OSPF Channel routing:** 20 points

The OSPF routing protocol should provide a list of nodes to which a channel message should be propagated. It must use shortest path finding and build the minimum spanning tree for each source node.

- **User Forwarding:** 30 points

Using the `PRIVMSG` command with a nickname target, the server should communicate with the daemon to get a next-hop from the local server to the server where the destination resides. Then, you must send a packet using a protocol of your devising. When an IRC server gets a forwarding packet, it should deliver the message locally or query the local daemon for the next hop and propagate the message. The message should travel along the path returned by the daemons and should ultimately be received. If path fails, you can drop the message and do not have to return an error.

- **Channel forwarding:** 15 points

Similarly, using the `PRIVMSG` command with a channel target, the server should communicate with the daemon to get a list of next-hops from the local server to the servers using that channel. Then, you must send a packet using a protocol of your devising. When an IRC server gets a forwarding packet, it should deliver the message locally and/or query the local daemon for the next list of hops given the source ID of the node from which it received the packet. If the path fails, you can drop the message and do not have to return an error.

- **Robustness:** 20 points

- Server robustness: 10 points
- Test cases: 10 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 a message longer than 512 bytes to test if there is a buffer overflow. We will make sure that your server does something reasonable when given an unknown command, or a command with invalid arguments. We will verify that your server correctly handles clients that leave abruptly (without sending a `QUIT` message). We

⁸ Just run 'irc' from the Andrew Linux command line.

will test that your server correctly handles concurrent requests from multiple clients, without blocking inappropriately.⁹

However, 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 20 points.

We will also look at your own documented test cases to evaluate how you tested your work.

- **Style: 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.

To help your development and testing, we suggest your server optionally take a verbosity level switch (-v level) as the command line argument to control how much information it will print. For example, -v 0 means nothing printed, -v 1 means basic logging of users signing on and off, -v 2 means logging every message event.

- **Extra credit: (up to 12 points)**
See extra credit section

13 Extra Credit

Our intent in suggesting extra credit items is to give interested students the opportunity to explore additional topics that we do not have time to cover in class. The primary reward for working on the suggested items is the additional experience and knowledge that they give you, not extra credit points. Extra credit will be granted at the discretion of the teaching staff.

For each suggestion, we list a rough estimate of the number of points you can receive. If you have more specific expectations about the extra credit you will receive, you should consult your TAs beforehand to avoid any disappointment.

If you work on the suggested topics below, please include in your project submission a file called `extra.txt`, describing what you have done. To receive credit for the programming items listed below, the code should be incorporated in your server. If it doesn't work, don't submit it! **Extra credit options should be disabled by default, but provide a command line argument to enable them.**

- **Test case, 3 points** – In general, your test code will be evaluated in the robustness part (see evaluation section). But you can get 3 points if your test code captures an interesting error case and is adopted for project grading.
- **Denial of Service, 6 points** – (3 points for implementation, 3 points for evaluation.) Section 8.10 of the RFC provides some suggestions for dealing with malicious clients that attempt to make the IRC service useless by flooding the network. Implement the suggestion, and experimentally evaluate its effectiveness. Report your experimental findings. Propose any other solutions to the problem that you think of.
- **Async name lookups, 3 points** – Your IRC server must perform DNS lookups on clients' addresses. As described in Section 8.11 of the RFC, the standard library calls may block for an extended time while the lookup times out. This is obviously an undesirable behavior for a server. Implement asynchronous name lookups using a separate thread to perform name lookups without blocking the

⁹ As an exception to this rule, your server may block while doing DNS lookups.

- primary server thread.
- **Scalability, 3 points** – Section 9 of the RFC notes that one of the current problems with the IRC protocol is scalability. Suggest how the protocol might be changed to improve scalability. You may get some ideas from looking at the design of IP multicast. (Or maybe not!)
- **Non-blocking Mini Protocol, 3 points** – An undesirable property of our routing daemon is that the IRC Server must block while waiting for a response from it. Make your IRC Server not block while waiting for a response. Furthermore, allow your IRC Server to submit multiple requests to the routing daemon at the same time; for example, so that it can initiate multiple route discoveries at the same time.
- **State synchronization (more like real OSPF), 10 points** – The link-state routing protocol we have implemented is extremely inefficient, because each node retransmits its entire state table every announcement period. Real OSPF implementations use high frequency “HELLO” packets to ensure that links are still alive, but only perform link state advertisements infrequently, or when a routing update occurs. Add this capability to your router – it will require modifying the packet format slightly to permit per-entry sequence numbers, and you’ll have to add new types to support the state synchronization and updates. **This will be a lot of work. 10 points is not that much.**

14 Getting Started

Depending on your previous experience, this project may be substantially larger than your previous programming projects. Expect the server implementation to require more than 1000 lines of code. With that in mind, this section gives suggestions for how to approach the project. Naturally, other approaches are possible, and you are free to use them.

- First, take a deep breath and do not panic.
- **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 your time to ask questions. For clarifications on this assignment, post to the main class bulletin board (academic.cs.15-441) 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.
- Decide how you will split up the work between you and your partner. It is recommended that both work together on the IRC server since it is an important part of the project and must be understood for the routing protocol. Some parts of this project can be done in parallel, but you should coordinate since they both have to work with the same program. Both of you should understand everything implemented for this project.
- Read the revised RFC 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 reading Sections 1-3. Do not focus on the details; just try to get a sense of how IRC works at a high level. Understand the role of the clients and the server. Understand what nicknames are, and how they are used. You may want to print the RFC, and mark it up to indicate which parts are important for this project, and which parts are not needed. You may need to reread these sections several times.
- Next, read Section 4 and 6 of the RFC. You will want to read them together. In general, Section 4

describes the purpose of the commands in the IRC protocol. But the details on the possible responses are given in Section 6. Again, do not focus on the details; just try to understand the commands at a high level. As before, you may want to mark up a printed copy to indicate which parts of the RFC are important for the project, and which parts are not needed.

- Now, go back and read Section 1-3 with an eye toward implementation. Mark the parts which contain details that you will need to write your server. Read project related parts in sections 4 and 6. Start thinking about the data structures your server will need to maintain. What information needs to be stored about each client?
- Get started with a simple server that accepts connections from multiple clients. It should take any message sent by any client, and “reflect” that message to all clients (including the sender of the message). This server will not be compatible with IRC clients, but the code you write for it *will* be useful for your final IRC server. Writing this simpler server will let you focus on the socket programming aspects of a server, without worrying about the details of the IRC protocol. Test this simple server with the simple IRC client `sircc`. A correct implementation of the simple server gives you 15 points for the core networking part.
- At this point, you are ready to write a standalone IRC server. Do not worry about the routing daemon, yet. This simple IRC server should run alone and implement all the features of the IRC protocol outlined above. In particular, it should allow users to log in, join channels, list channels, and send messages to either users or channels. This will go a very good start towards your final IRC server and will give you an idea of how your final IRC server should work. Make sure this server is solid. You will be building upon it and you do not want to waste time trying to figure out if a particular error is in the IRC server or the routing daemon. **Remember only to use ports in the range of your group number to avoid weird conflicts!**
- Do not try to write the whole server at once. Decompose the problem so that each piece is manageable and *testable*. Read related parts of RFC again carefully and think about how the commands work together. For each command, 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 the simple client `sircc` or `telnet`.
- Thoroughly test the IRC server. Use the provided scripts to test basic functionality. For further testing, use the provided `sircc` client or `telnet`. It may be useful to learn the basics of a scripting language to make some repeatable “regression tests.” As said, the routing daemon will build upon the IRC server, so thorough testing of the latter will save time in debugging the former.
- Once the simple IRC server is complete, you can begin worrying about the routing daemon. Since the routing daemon and IRC server are two programs, it is not necessary to have the two talking to each other yet. In fact, it is probably best to have the two thoroughly tested independently to trap errors more efficiently. First, get familiar with UDP socket programming, which is almost identical to TCP socket programming, but not quite. There are some references to help you at the end of this document. Next, write up a design of each part of the routing daemon and decide what data structures you will need.
- First tackle general flooding and table construction. Work on getting the link entry and user entry tables functional. Once the protocol works for messages between users, then start working on multicasting, minimum spanning trees, shortest path finding, and multicasting.
- Again, thoroughly test the routing daemon. `Telnet` is a very useful tool for this. Make you’re your daemon can add paths, remove paths, find paths, withstand failures, and does not segfault.
- Before you start implementing message forwarding in your IRC Server, carefully design a

protocol. You might need to differentiate between users and channels since channels need to be multicasted. Outline what parts of the original IRC server need to be modified in order to connect and talk to the routing daemon.

- Almost there! Hopefully, after implementing the message forwarding protocol and server extensions everything will work perfectly. More likely, though, something will break. Things that work perfectly separately do not always work perfectly together. This is a big software engineering problem. So, yet again, thoroughly test the final product. Run the same tests you used on the individual pieces to make sure nothing broke in the merge.
- You may use some of the system call wrappers provided by CS 15-213 `csapp` library (included with the simple IRC client package). However, for server robustness, you should not use certain wrappers such as `Select` since temporary system call failures (e.g., `EINTR`) would cause the server to abort. Instead, your server should handle such errors gracefully. For the same reason, you should NOT use the `RIO` read/write functions provided by the `csapp` library as they may cause your server to block while reading/writing, or give inappropriate return codes.
- “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!**

15 Resources

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

- Class Textbook – Sockets, OSPF, etc
- Class B-board – Announcements, clarifications, etc
- Class Website – Announcements, errata, etc
- Computer Systems: A Programmer’s Perspective (CS 15-213 text book)¹¹
- BSD Sockets: A Quick And Dirty Primer¹²
- An Introductory 4.4 BSD Interprocess Communication Tutorial¹³
- Unix Socket FAQ¹⁴
- UDP Sockets Slides¹⁵
- TCP Sockets Slides¹⁶
- Sockets section of the GNU C Library manual
 - Installed locally: `info libc`
 - Available online: GNU C Library manual¹⁷
- man pages
 - Installed locally (e.g. `man socket`)

¹⁰ <http://www.ietf.org/rfc/rfc1122.txt>, page 11

¹¹ <http://csapp.cs.cmu.edu>

¹² <http://world.std.com/~jimf/papers/sockets/sockets.html>

¹³ <http://docs.freebsd.org/44doc/psd/20.ipctut/paper.pdf>

¹⁴ <http://www.developerweb.net/forum/viewforum.php?f=53>

¹⁵ http://www.cs.rpi.edu/courses/netprog/lectures/pphtml/udp_sockets_programming/

¹⁶ http://www.cs.rpi.edu/courses/netprog/lectures/pphtml/tcp_sockets/

¹⁷ http://www.gnu.org/manual/glibc-2.2.3/html_node/libc_toc.html

- Available online: the Single Unix Specification¹⁸
- `select()` tutorial¹⁹
- Google groups²⁰ - Answers to almost anything

¹⁸ <http://www.opengroup.org/onlinepubs/007908799/>

¹⁹ http://www.phildowd.com:4075/phpMan.php?docType=man&parm=2%20select_tut

²⁰ <http://groups.google.com>