

# IP for Minet

## Project Part C

**OPTIONAL FOR FALL 2000 - WILL NOT COUNT FOR YOUR GRADE**

### Overview

In this part of the project, you and your partner will build an implementation of IP for the Minet TCP/IP stack. You will be provided with all parts of the stack except for the IP module. You may implement the module in any way that you wish so long as it conforms to its interfaces to other Minet components, and to the reduced IP standard described here. However, Minet provides a considerable amount of code, in the form of C++ classes, that you may use in your implementation. Furthermore, we recommend that you implement your module as a select-based server, similar in spirit to the third part of project part A. You may also earn extra credit by implementing additional parts of the IP standard.

### The Minet TCP/IP Stack

The Minet TCP/IP Stack is documented in a separate handout. After completing project part B, you should be quite familiar with Minet. You will be given the source code to all of the Minet modules, except for `ip_module`. The `ip_module` source you will be given is only a framework – it simply connects with the rest of the stack. You are responsible for filling it out.

You will also receive binaries for reader, and writer. It is vital that you use these programs because they must be run as root or else they will not work. Because we don't want to give you root access, we are giving you binaries that are setuid to root. This means that these programs (and only these programs) run as root. If you would like to use reader and writer on a machine outside the TLAB, you will need root privileges on that machine.

### Your IP Addresses

Each of you was assigned 255 IP addresses (a 10.10.x/24 network address) earlier in the quarter. You must again use these addresses for this project. They are special in that packets sent to them will not be forwarded beyond the local network. In fact, if you are using machines other than the TLAB machines, you will need to add a route so that they actually make it to the local network.

### Dedicated TLAB Machines

You may use any of the TLAB machines, either from the console or remotely via ssh. In fact, you'll usually want to use two of them simultaneously. We have dedicated TLAB-11 through TLAB-15 to running Linux for the duration of the quarter. These machines should always be available for remote or console login. If they are not, send mail to [request@cs.cmu.edu](mailto:request@cs.cmu.edu).

## IP Specification

The core specification for IP is RFC 791, which you can and should fetch from [www.ietf.org](http://www.ietf.org). In general, you will implement IP as defined in that document, except for the parts listed below.

- You do not have to support the Type Of Service field, other than copying it. Your service may be completely FIFO.
- You do not have to implement any IP options.
- You do not have to support IP fragmentation. You should limit yourself to sending 576 byte packets and always set the Don't Fragment flag. This size IP packet is guaranteed not to be fragmented by the network, and all IP implementations must support at least this size packet.
- You may assume that the MTU is 1500 (ie, that everything is an Ethernet).
- You do not have to send or handle ICMP packets.
- You must implement routing and support at least two interfaces. You must initialize your routing table from the file `“./minet_static_routes”`

Chapter 4 of your textbook covers IP within the context of routing. We suggest that you read it first, then read RFC 1180, and finally read RFC 791. In addition, we will be handing out other information as the quarter progresses. Rick Steven's book, "TCP/IP Illustrated, Volume 1: The Protocols" is also very useful, as is Doug Comer's "Internetworking With TCP/IP Volume I: Principles, Protocols, and Architecture". Both of these books are available from Peter.

## Interfaces

You must fully support the following two interfaces.

- `eth0` : The first Ethernet card
- `lo` : The local loopback device

In addition, you must implement support for an arbitrary number of Ethernet devices (`eth1`, `eth2`, etc). The local loopback device forwards IP packets right back up the stack. It is used to allow two processes that are sharing the same Minet stack to communicate using IP packets. The IP address of `lo` is 127.0.0.1. `eth0` will be the destination for most packets bound for other machines. You must assign it the IP address in the environment variable `MINET_IPADDR`. Its physical address is in `MINET_ETHERNETADDR`. The addresses of all Ethernet cards must be 10-net addresses.

## Routing

You must implement support for routing over all interfaces. This means the following:

- Locally generated packets bound for local addresses (ie, the address of any of your interfaces) should be routed to `lo` so that they go back up the stack.
- Locally generated packets bound for remote addresses should be routed through the most appropriate interface according to the longest prefix matching rule.

- Remotely generated packets bound for local addresses should be routed up the stack.
- Remotely generated packets bound for remote addresses that are routable with your routing table should be routed out the appropriate interface according to the longest prefix matching rule.
- Remotely generated packets bound for remote addresses that are not routable should be dropped.

### **IP Module / IP Multiplexor Interface**

The IP module talks to the IP multiplexor using Packets that have extracted IP Headers. The Packets may have other headers, such as Ethernet, UDP, TCP, or ICMP headers, however, an IP header must always be present. When a received IP packet is bound for a local address, it should be serialized to the IP multiplexor. Similarly, the IP multiplexor will serialize outgoing IP packets to the IP module.

### **IP Module / Ethernet Multiplexor Interface**

The IP module communicates with the Ethernet multiplexor using RawEthernetPackets. To send a packet, it serializes a RawEthernetPacket to the multiplexor. To receive a packet, it unserializes a RawEthernetPacket from the multiplexor.

### **IP Module / ARP Module Interface**

The interface between these modules is request response. To make an ARP request, the IP module serializes an ARPRequestResponse with the flag REQUEST to the ARP module. In response, the ARP module will serialize back an ARPRequestResponse with the flag RESPONSE\_OK (meaning that the IP address was found in the cache and that the mapping in the response is valid), or RESPONSE\_UNKNOWN (meaning that no mapping was in the cache and thus the response's mapping is invalid.) Generally, when a RESPONSE\_UNKNOWN is returned, the ARP module will have also made a broadcast ARP request. The IP module may retry the request at any time.

### **Recommended Approach**

There are many ways you can approach this project. The only requirements are that you meet the IP specification detailed above, that your IP module interfaces correctly to the rest of the Minet stack, and that your code builds and works on the TLAB machines. We recommend, however, that you use C++ and exploit the various classes and source code available in the Minet TCP/IP stack. Furthermore, we recommend you take the roughly the following approach.

1. Read Chapter four of your textbook
2. Read RFC 1180
3. Read RFC 791.
4. Reread the "Minet TCP/IP Stack" handout.
5. Fetch, configure, and build Minet. We will provide an `ip_module.cc` that consists of the framework code needed to start up and connect to the appropriate fifos.

6. Extend this IP module so that it simply prints incoming RawEthernetPackets from the Ethernet multiplexor and locally generated Packets from the IP multiplexor. You should be able to run the stack with `udp_client` or `tcp_client` as the application and see packets printed when you type and when you send your stack traffic from netcat (`nc`).
7. Write a simple `lo` interface and router that forwards IP packets bound for 127.0.0.1 back up the stack. You should be able to test this using `udp_client`.
8. Write a generic interface class and specialize it for Ethernet. This class will store state about the interface such as: (0) it's name, (1) whether it is up or down, (2) its IP address, (3) its physical layer (Ethernet) address, and (4) which fifos to use to talk to it.
9. Write a routing table class. A routing table consists of mappings of network addresses (i.e., 10.10.0.0/16 or 10.10.0.0 netmask 255.255.0.0) to interfaces. Your routing table should be able to initialize itself with the mappings stored in a file (the `minet_static_routes` file, which we will give you). Your routing table should support dynamic changes (like those induced by RIP messages).
10. Implement the longest matching prefix rule within your routing table class. When given an IP address, your routing table should be able to respond with an interface.
11. Add static routes that send 10.10.0.0/16 packets to `eth0` and that sends packets bound for local IP addresses up the stack.
12. Add code to construct an IP Packet from each RawEthernetPacket you receive from the Ethernet multiplexor.
13. Add code so that each Packet you construct from an incoming Ethernet packet is routed according to your routing table.
14. Test incoming packets using netcat in UDP mode.
15. Add code so that each Packet you receive from the IP multiplexor is routed using your routing table.
16. Interface with the ARP module so that you can find hardware addresses for these outgoing packets. You'll need to put these destination addresses into an Ethernet header and prepend it to the outgoing packet.
17. Add code to construct a RawEthernetPacket from such an outgoing packet and to serialize it out to the Ethernet multiplexor.
18. Test outgoing packets using `udp_client`.
19. Test bi-directional communication using `tcp_client`.
20. Test your routing table by adding clones of your `eth0` interface bound to different addresses.
21. You're done!

### **Extra Credit: IP Fragmentation Support**

For extra credit, you may implement support for IP fragmentation. If you do so, you should implement support both for fragmenting packets and for reassembling them.

## Extra Credit: RIP Support

For extra credit, you may implement support for the RIP routing information protocol, so that you can advertise your routes and take advantage of the route advertisements of others. .

## Caveat Emptor

The Minet TCP/IP Stack is a work in progress. You are the first class to use it. You can and will find bugs in it. We will do our best to make your experience as pleasant and fair as possible. If you try hard, you will do well. If some part of Minet turns out to be an unexpected impediment, this will not be reflected in your grade.

We also appreciate constructive feedback and suggestions. We plan to use Minet extensively for teaching in the future. Your ideas, which we will credit accordingly, can therefore have a lot of impact.

## Mechanics

- Your code must function as a `ip_module` within the Minet TCP/IP Stack, as described in a separate handout.
- Your code should be written in C or C++ and must compile and run under Red Hat Linux 6.2 on the machines in the TLAB. In particular, we will compile your code using GCC 2.95.2 and GNU Make 3.78.1, which are installed on the lab machines. You must provide a Makefile. We will expect that running “make” will generate the executable `ip_module` and that this module will meet the specification described in this document and in the “Minet TCP/IP Stack” handout.

## Things That May Help You

- RFC 791 is essential.
- RFC 1180 is a handy tutorial on IP and the lower layers it talks to.
- Chapter 4 of your book. Section 4.4 is a reasonable introduction. Remember that you don't have to implement ICMP, routing protocols, or fragmentation.
- Rick Stevens, “TCP/IP Illustrated, Volume 1: The Protocols”
- Doug Comer, “Internetworking With TCP/IP Volume I: Principles, Protocols, and Architecture”
- Rick Steven, “Advanced Programming in the Unix Environment” – especially handy if you are having issues with fifos.
- The handout “Unix Systems Programming in a Nutshell”
- The handout “Make in a Nutshell”
- The handout “The TLAB Cluster”
- The C++ Standard Template Library. Herb Schildt's “STL Programming From the Ground Up” appears to be a good introduction
- GDB, Xemacs, CVS, etc.