Internet Protocol Review/Introduction

This lecture introduces the the most common network protocol, Internet Protocol (IP) version 4 (IPv4).
After this lecture you should be able to: differentiate between the Internet and IP; look up IP standards; interpret the values of the most common IP header fields; determine the netmask for an IP network; determine if an IP address is in a particular network; determine if an IP address is public, private or link-local; decide which port a packet would be forwarded on based on the contents of a routing table.

Introduction

IP was developed in the late 70's when different computer manufacturers, academic institutions and research groups used incompatible data communication networks and protocols. IP was designed as a common protocol to link these networks together so they could exchange files, e-mail, terminal sessions, etc. It was thus an inter-network protocol or an "internet protocol".

As the usefulness of a universal networking protocol became clear, new system started using IP as their *native* networking protocol. Widespread adoption of IP has resulted in almost all networks using IP as their network-layer protocol.

The availability of IP, a widely-supported and freely-available protocol, facilitated the growth of a non-proprietary commercial data network using IP that is commonly called "The Internet".

Exercise 1: What is the difference between IP and "The Internet"? Does a network using IP have to be on the Internet? Does someone using the Internet have to use IP?

IP is defined in documents called Requests For Comment (RFCs) published by the Internet Engineering Task Force (IETF). IETF standards development is open to the public and adoption of proposals depends on technical merit.

IP is a network (Layer 3) protocol and handles the routing of packets between (potentially) different local networks. The geographic scope of these networks can be very large (global).

IPv4

The current version of IP is known as IP version 4 (IPv4) and is defined in RFC 791, published in 1981.

IP Version 6 (IPv6) was originally designed to address the problem of address space exhaustion but introduced additional, unrelated changes. Simpler methods were developed to work around the address exhaustion problem and it remains to be seen if IPv6 will ever be widely used.

Not surprisingly, documentation for IP protocols is widely available on the Internet, for example, from www.ietf.org.

IP is a packet-switching protocol. Data is encapsulated in IP packets which are transferred in store-and-forward manner between routers from a source host to a destination host:

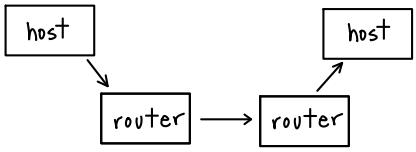

IP Packet Format

Services provided by the IP layer are limited to routing and fragmentation. IP has "... no mechanisms to augment data reliability, flow control, sequencing, ..."[1].

The IP header consists of a minimum of five 32-bit words (20 bytes). The following diagram is taken from RFC791[2] and shows the IP header:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Version|  IHL  |Type of Service|          Total Length         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Identification        |Flags|      Fragment Offset    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Time to Live |    Protocol   |         Header Checksum       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Source Address                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Destination Address                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Options                    |    Padding    |
```

[1] John Postel, Ed., RFC 791.
[2] RFCs were published in text format.

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Example Internet Datagram Header

Figure 4.

The most important fields are:

Version Protocol version number (4)

IHL IP Header Length in 32-bit words

Type of Service Priority. Not widely used.

Total Length Length of the IP packet in bytes.

Identification/Flags/Fragment Offset for fragmentation. Rarely used.

Time to Live A value that is decremented each time a packet is forwarded. Prevents packets traversing routing loops indefinitely.

Protocol the type of protocol embedded in the IP frame. 1 for ICMP, 6 for TCP, 17 for UDP. Assigned by IANA[3].

Header Checksum a one's-complement checksum for the header

Source/Destination Address the 32-bit source and destination IP addresses (see below).

Options Optional header components that are not normally used (security, source routing, route recording and timestamps).

IPv4 Addresses

Host Addresses

Each IP (IPv4) network interface has a 32-bit (4 byte) address. Most hosts, and all routers, have more than one interface and thus more than one IP address.

For example, a typical home router will have two IP interfaces, one public and one private, as well as a software-based loopback interface:

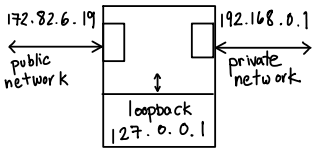

IP addresses are usually written as a "dotted quad" of the decimal value of each byte separated by periods. The bytes are written in big-endian order. For example, 0xc0a80001 would be written 192.168.0.1.

[3] Internet Assigned Numbers Authority.

Most devices with IP protocol stacks have a virtual network interface at address 127.0.0.1 (hostname localhost) which is used for communication between processes on the same device.

Network Addressses

IP addresses are assigned in a hierarchical manner. The most significant bits of the address identify one of a few hundred thousand (up to perhaps 10^6) IP networks in the world. The remaining bits identify a host within that network.

Originally networks were divided up into three classes: A, B and C. Class A networks could have up to 2^{24} host addresses. Class B addresses up to 2^{16} and Class C up to 2^8.

This led to inefficient allocation of network addresses and today network address are "classless" and are composed of two parts: the value of the network prefix (e.g. 142.232.0.0) and the length of the network portion of the address in bits preceded by a slash (e.g. /16). The two values together are the (classless) network address. For example, the BCITNET2 network has an address of 142.232.0.0/16.

A netmask is a 32-bit value with 1's in the bits corresponding to the network address.

Exercise 2: What is the netmask in binary for a /24 network? What is it in decimal? How can the netmask be used to determine if one IP address is on the same network as another? Is the address 192.168.2.200 in the 192.168.2.0/25 network?

Network addresses are assigned by a non-profit organization called ICANN. The host whois.arin.net can be used to query for ownership of North American network addresses.

Exercise 3: Who "owns" the 204.191.0.0/16 network?

IP Routing

IP networks operate in store-and-forward fashion. Routing is the process of getting a packet from source to destination.

IP networks use routers connected in a mesh which can, and often does, contain redundant links and loops.

Each IP packet includes a "time to live" field to protect against packets circulating in the network indefinitely as a result of misconfigured routers.

Each packet is routed independently. Thus each packet has to have the destination address.

Each device that forwards IP packets typically has multiple ports and is called a "router" because it decide on which port(s) to forward the packet. This decision is done by looking up the destination IP address in a "routing" table that defines the outgoing port for a network.

Each routing table entry can have a cost or "metric" associated with it. The routing algorithm selects the lowest cost route (port). Routing costs are local to each router. The metrics can be determined in many ways. For a simple host they may be based on the port data rate. Some routers exchange information with other routers to determine the best route to each network. The metrics can also be modified dynamically by factors such as delay. Routes can also be manually configured.

A routing table will usually have a default route containing the address of a "gateway" router. Packets for which there is no route are send with a destination L2 address of the gateway. The gateway receives these packets forwards them based on its own routing table.

Here is the routing table for a simple home wireless router:

Destination	Gateway	Subnet Mask	Metric	Interface
192.168.1.0	*	255.255.255.0	0	br0 (LAN)
204.191.0.0	*	255.255.0.0	0	vlan1 (WAN)
127.0.0.0	*	255.0.0.0	0	lo
default	204.191.1.1	0.0.0.0	0	vlan1 (WAN)

Exercise 4: For the routing table above, what port ("Interface") would be used by packets with the following destination IP addresses: 127.0.0.255? 192.168.1.1? 192.168.2.1? 204.191.10.32?

Classes of Routers

Many routers like the one above contain only a few routes, often just for the attached network segment and a default route pointing to an "upstream" router.

However, routers that connect multiple networks need to determine which port should be used to reach different networks. This can be done manually or using distributed algorithms that try to determine the best route from each router to each network. The most common of these algorithms is called OSPF (Open Shortest Path First). Most "enterprise" routers will configure their routes using OSPF.

Routers that connect different service providers are called "border" routers. Their routing decisions are more complex because they need to take into account "peering" agreements between service providers that determine which packets can be forwarded to which service providers. The protocol typically used to set up routing between border routers is BGP (Border Gateway Protocol). Border routers are typically found only at a one (or a few) data centers in each city called IXP (Inter-Exchange Points) where ISP interconnect their networks.

In future lectures we will study the operation of these routing protocols in detail.