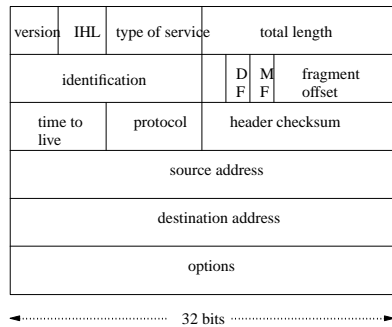


IP fields

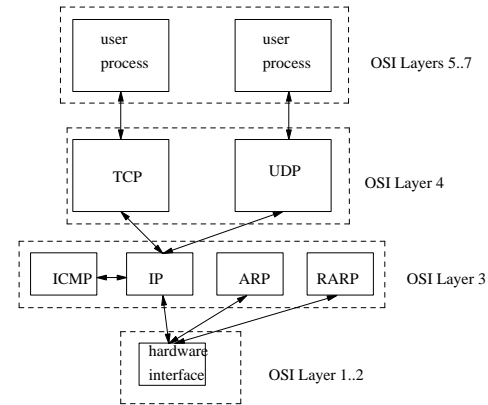


IP fields

- total length
 - length of both header and data
- identification
 - all fragments of a datagram contain the same id value host can determine which datagram an incoming fragment belongs
- DF
 - do not fragment
- more fragments
 - All fragments except last one must have this bit set to true

TCP/IP Support Protocols

- are another reason TCP/IPs popularity



IP datagram structure

- an IP datagram consists of a header parts and a text part
 - header has a 20 byte fixed part and a variable length optional part
- *type of service* field allows different combinations of reliability and speed to be chosen
 - for digital speech IP can be told to emphasize fast delivery
 - for file transfer is taking IP can be told that accuracy is paramount at the expense of speed

IP addressing

- *internet is a virtual structure*
 - implemented entirely in software
 - packet frames and addresses were designed on merit

- addresses contained with 4 bytes
 - conceptually the 32 bit number has two parts
 - hostid
 - netid

- three primary classes of IP addresses

IP fields

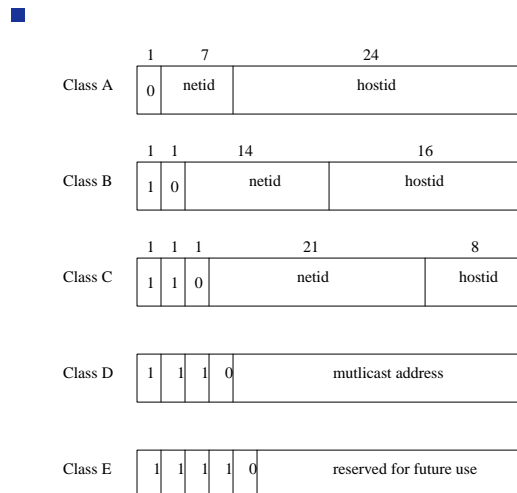
- Fragment offset
 - must be a multiple of 8. Tells receiver where this fragment belongs in the datagram.

- Time to live
 - in seconds. Decrements each second or each hop, when it reaches 0 it is thrown away

- Protocol field
 - tells which of the various transport processes the datagram belongs. Ie TCP or UDP

- Header checksum
 - verifies header only

IP address classes



IP fields

- Source and destination address
 - indicate the network number

Network structure

- machine R is a router
 - has at least two IP numbers - one for each network card

- *because Internet addresses encode both a network and a host on that network; they do not specify a host, but a connection to a network*

Address Classes

- class A
 - a handful of network which have more than 65536 hosts

- class B
 - addresses for intermediate size networks. 256..65535 hosts. 14 bits for netid 16 bits for hostid

- class C
 - networks which have less than 256 hosts

- class D
 - multicast, hosts may dynamically join/leave multicast group
 - hosts may be in many different multicast groups

Network and Broadcast addresses

- two reserved hostid's


- Internet addresses can be used to refer to networks as well as network cards. By convention the network address has hostid all bits 0

- a broadcast address conversely has hostid bits all 1

- one of the weakness of IP addressing is that if a machine changes network - its IP address must change

Network structure

- consider the following IP network



IP Support Protocols

- ICMP (Internet Control Message Protocol)
 - sends control information between the hosts
 - routers generate most of this information

- routers use ICMP to
 - inform hosts that a packet could not be delivered because of an error
 - or a better route exists to a particular destination

- ICMP messages are send using IP frames
 - ICMP messages use the IP *protocol* field and set it to 1

Typical ICMP messages are

- *destination unreachable* - when a router cannot find a routing table entry for the destination of an IP packet

- *routing redirect* - a router sends a routing redirect message to inform a host that a better route exists via another router

- *time expired* - message indicates a packets *ttl* field has reached 0
 - usually because of a configuration error
 - malfunctioning router

- *echo request and echo reply* - echo request messages request that the destination return the data in an echo reply message (ping)

Fragmentation and reassembly

- IP datagrams may be fragmented *en route*
 - if intermediate nodes cannot cope with a large datagram (MTU (maximum transmission unit) is smaller than datagram size)

- IP datagrams may be reassembled *en route*
 - although not a good idea as routing is dynamic. (So datagrams may not always travel the same route)

- to fragment a datagram into two a node creates two new datagrams with same fragment ids
 - the first offset is 0, MF = 1
 - the second offset is n, MF = 0

Fragmentation and reassembly

- consider trying to send a 1420 byte datagram when the MTU is 620
 - $1420 = 1400 \text{ data} + 20 \text{ IP header}$

- split into 3 packets
 - first packet length = 620 = 20 new IP header + 600 old data, offset 0
 - second packet length = 620 = 20 new IP header + 600 old data, offset 600
 - third packet length = 220 = 20 new IP header + 200 old data, offset 1200
 - the new fragments have the same unique frag id as the original why?

- reassembly reverses this process

RARP (Reverse ARP)

- ARP maps from network addresses to datalink addresses
 - sometimes you require the opposite mapping
- many machines can read their datalink hardware to find out the datalink address
 - but then needs to find out its IP address
 - for example, disk less workstation, X terminal, printer

RARP (Reverse ARP)

- at least one host on the IP network must contain a list of IP addresses with corresponding datalink addresses
 - whereas ARP does not require that this list is present
 - a RARP is a broadcast request - any host may reply
- the sender fills in its datalink address
 - its network address is filled with zeros
 - specifies the target datalink address (normally the same as sender)
 - the RARP server fills in the requested IP (network) address

Address Resolution Protocol

- IP address space is virtual and has no addressing relationship with the underlying datalink protocols
 - every network interface has an IP address
 - every network interface has a datalink address
 - datalink addresses vary in format and size
- suppose IP is sending a packet to a remote host on the same Ethernet
 - IP needs destination Ethernet address
 - could manually keep track of hosts and their interface card datalink addresses

Address Resolution Protocol

- clearly on a large network this becomes unmanageable
 - ARP (Address Resolution Protocol) is an automatic method which maps any network level address (IP address) to datalink address
 - ARP does this by exploiting the broadcast capability commonly found in most LAN datalink protocols

TCP and UDP

- primarily there are two transport protocols used with IP: TCP and UDP
 - remember that IP may provide an unreliable service
- **Transmission Control Protocol (TCP)**
 - provides a flexible two-way byte stream protocol (byte stream allows addressing *within* a host - to user, process or service)
 - *provides a bidirectional pipe*
 - the source and destination address are called a *Port*
 - TCP is the most popular transport protocol on top of IP
 - it uses sliding window technique to provide a reliable service
 - it uses a three way handshake to establish a connection
 - and a two way handshake to disconnect

RARP (cont)


- RARP is normally the first step taken when a diskless workstation is powered up
 - once it knows its own IP address it can then proceed to load its operating system from a network server by using a simple file transfer protocol (TFTP)

TCP

- typically it is used for:
 - mail
 - file transfer
 - remote login protocols
- technically it provides
 - a reliable duplex byte stream
 - eliminating duplicate packets
 - handling retransmission of lost packets
 - and ensuring data is delivered in order

TCP/IP services

- Internet services are usually designed and implemented using the client/server model of computing
- divides the service into three distinct parts.



- each performs a distinct function in implementing the client/server model:
 - client process. The user who is using this service
 - server process. Maybe on another machine - providing the service
 - the protocol which the client and server are using to communicate

TCP

- **User Datagram Protocol (UDP)**
 - is an unreliable datagram protocol and is deliberately simple
 - it does not ensure that packets arrive in order, un duplicated, or even at all!
 - it sends discrete datagrams, and delivers messages that arrive to the appropriate *Port* (same addressing schema as TCP)
 - a *port* may belong to a user, process or service
 - the standard Internet name service, DNS, uses UDP
 - it can be regarded as multiplexing many users, processes and services through one IP address
 - UDP has no standard connection procedure and no disconnect procedure

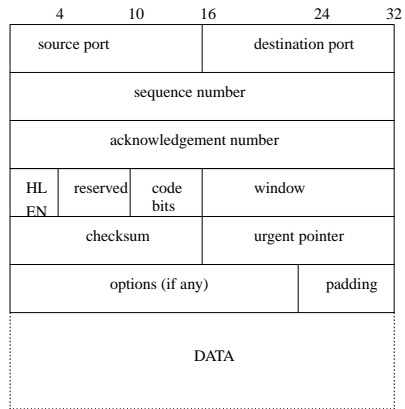
References

- refer the interested reader to explore UNIX BSD socket interface

Stevens *UNIX Network Programming*, Prentice-Hall 1990

Bach *The design of the UNIX operating system*, Prentice-Hall 1986

TCP Header format



TCP Header format


- HLEN 4 bits
 - header length = $n \times 32$

CODE field meaning	
URG	urgent pointer field valid
ACK	ack field valid
PSH	segment needs push
RST	reset connection
SYN	synchronize sequence no.s
FIN	sender reached end

TCP/IP services

- note that this client server model is different to that discussed by business schools and management information systems as the client server model used by TCP/IP involves processes running on machines which are *not* dedicated as clients or servers
 - using TCP/IP users may dictate which machines are client or servers
 - one machine maybe a client for one application
 - and a server for another
- TCP/IP clients and servers are software processes
- whereas MIS treats clients and servers as machines (Ie mainframes and PCs)

TCP/UDP model



Telnet

- uses TCP
- Virtual Terminal Protocol
 - was the Internet standard for remote login (now replaced by ssh)
- users can remotely login to another machine on the same LAN or any machine on the TCP/IP offering this service
 - eg. login to a machine in Japan
 - passwords are sent in plain text

ssh

- uses TCP
- a secure replacement for ftp and telnet
- both ftp and telnet should be retired if at all possible as they pass login name and password across the network in plain text..
- ssh also comes with scp
 - secure shell copy
- sftp, secure shell ftp
 - all username passwords and data contents are heavily encrypted

Services offered by TCP/IP

■

Layer	Protocols												
Process Application	NN TP	HT TP	PO P3	RC MDS	X rpc xdr	NFS	D N S	TF TP	SN NP	N T p	DH CP		
Transport	TCP						UDP						
Internet	[ICMP, ARP, RARP]						IP [EGP, BGP, IGMP]						
Network	Ethernet			X25			SLIP			PPP			

File Transfer Protocol

- uses TCP
- allows users to transfer files from one machine to another machine providing that the user has an account on both machines
- requires little configuration
- a special FTP service called anonymous FTP does not require users to have accounts on both machines but needs more configuration

Services (continued)

- **Printer Spooling**
 - many PC users still think erroneously that print spooling is a network's only use
 - the most common UNIX print spooling protocol is **Line Printing Protocol (LPP)**" this was introduced in 4.2 BSD (released in 1981)
- allows a central server to be connected to a print resource
 - users can access this resource for other clients

Services offered by TCP/IP

- **inetd: the servers server**
 - a host wanting to support many services could require a process for a server for each service
 - even on UNIX systems, the number of idle processes can quickly grow beyond a reasonable size, thus absorbing system resources unnecessarily (servers typically spend most of their time waiting for connections from clients)
 - **inetd** server accepts connections for a range of protocols and invokes the appropriate servers when needed

Services (continued)

- **System logging (syslog)**
 - important for system administrators - it permits logging information to be directed to any host on the network.
- **Remote Shell Protocol (RSH)**
 - permits the user to run a shell command on the remote machine. (A form of remote job execution). It uses the same protocol as **rcp** (remote copy) and **rlogin** (remote login)
 - commonly known as the R commands

Services (continued)

- **Network File System (NFS)**
 - provides transparent file access over the network so that remote file appear as if they are local
 - NFS was designed for use over the LANs and is widely used in such environments
- NFS is widely available for different operating systems: UNIX, Macintosh, DOS clients exist.
 - (Ie at University. of Glamorgan. we use NFS between GNU/Linux, Sun OS)

Naming and infrastructure

- there is a need for a mapping of textual domain names to numeric IP addresses
 - difficult to remember 193.63.130.52 is the class C address for floppsie!
- also require services such as a consistent time between different machines

A few R commands

- **rup**
 - list all machines which are up and running and report their load average
- **rusers**
 - list all users logged in over the (local!) network
- **rwho**
 - list all (local) users and where they are logged in

Domain Name Service (DNS)

- the Internet standard DNS maps host names, such as `floppsie.comp.glam.ac.uk` to IP addresses such as `193.63.130.52`
- DNS namespace is partitioned hierarchically into a tree
 - `glam.ac.uk` - may map onto several class C networks
 - `floppsie.comp` - indicates a machine within the computer studies network
 - an interface card on class C network `193.63.130.xx`

More R commands

- **netfind**
 - try to find a particular user and give this persons email address.
- **rep**
 - allows a user to copy files between machines ie:
`rcp merlin:/user/fred/foo`
`lancelot:/user/mary/bar`
 copying from source to destination

Communication services

- many different services provided above TCP/IP
 - one of the most widespread use of computer networks is to communicate with other people

- one of the oldest utility is **Electronic mail**
 - users compose a message with a **user agent (UA)**
 - a user agent is an editor which passes the composed message to a MTA
 - one of the oldest is the UNIX program /bin/mail
 - one of the newest is outlook

Naming and infrastructure (continued)

- we could have a simple lookup table that is manually updated
 - soon becomes unmanageable
 - use a dynamic mechanism, domain name service

- have a machine which will keep track of IP addresses and ASCII names
 - if it cannot resolve a name it requests help from another machine higher up the tree
 - the DNS protocol specifies how DNS clients ask DNS servers for mappings
 - and how DNS servers communicate with each other.

Communication services

- these programs send the message to a **message transfer agent (MTA)**
 - typically this information is sent using SMTP (**Simple Mail Transfer Protocol**)
 - SMTP is the Internet standard protocol and all MTAs on the Internet use SMTP.
 - one of the dangers of SMTP is that it is relatively insecure (it is possible to forge email messages!)

- **Network News Transfer Protocol (NNTP)** protocol for delivering and accessing USENET news over Internet

Network Time Protocol

- used so that different hosts can keep the same time of day

- required by many applications - email to ensure that you don't receive a message before it was sent!

IP Configuration

- *every* IP address actually refers to the interface card and **NOT** the machine!
- thus a gateway machine will have at least two interface cards
- to add a new machine floppsie onto the Computer Studies network
 - floppsie's interface card has to be assigned a unique IP address
 - first three numbers the same as the Computer Studies network (193.63.130)
 - class C network - means first 3 bytes are always the same
 - last number is the interface card number (hostid) 52
 - 193.63.130.52

IP Configuration

- software on the new machine needs to know:
 - the gateway on the Computer Studies network to other networks
 - its own interface card IP address!
 - the *nameserver* IP address. The *nameserver* translates all ASCII names to IP numerical addresses.
 - hop metric given with each gateway

Example IP configuration

- Case study - adding a machine onto the Computer Studies network
- the IP network in Computer Studies connects:
 - Apple computers
 - NT
 - SPARC machines
 - GNU/Linux
 - routers and various print services, etc

Example IP configuration

- each machine may run a different protocol above the IP layer if they wish
 - but most machines will run the IP protocol
- the Computer Studies IP network is connected (via a gateway) to the University of Glamorgan IP network
 - in turn is connected to the world IP network via another gateway (through the University of Glamorgan 1M bit line)

IP Configuration

