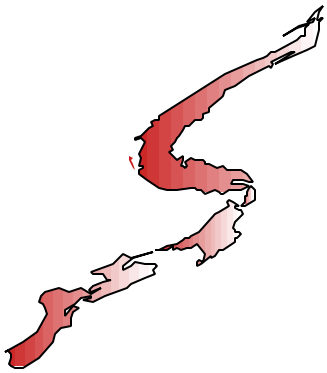

Network Layer continued

Dr. G. A. Marin

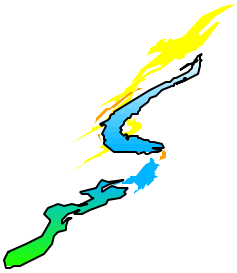

IP Address

- Each host and router on the Internet has an unique IP address.
- Address is 32 bits long
- Written as four decimal numbers d1.d2.d3.d4
- Four formats have been defined for IP addresses:
  - Class A: allows 126 netid's and 16,777,216 hostid's
  - Class B: 16,382 netid's and 64K hostid's
  - Class C: 2M netid's and 254 hostid's
  - Class D: Multicast addresses
- Network numbers uniquely assigned by Network Information Center
- Address 0.0.0.0 used only by host when booting.
- All 1's means broadcast on local net.
- Correct netid plus 1's means broadcast to netid.

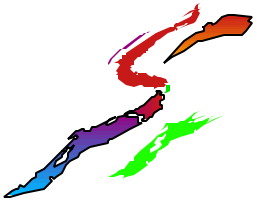

Address Formats by Class

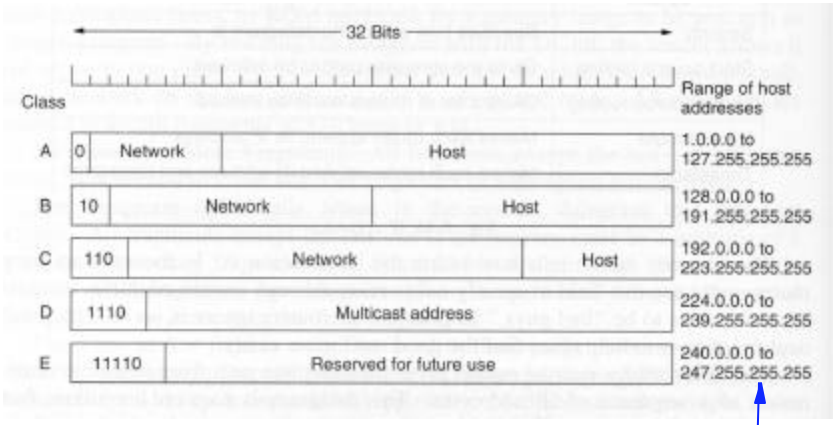

Illustrates numerical range and NOT interpretation.

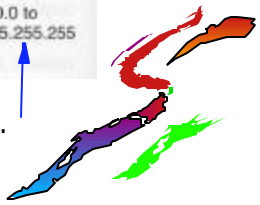

Address Administration Problems

- Every time a new network is installed the system administrator must contact NIC to get a new netid?
- Moving a machine to a new LAN requires a new address, changes to config files, advertising the address...
- If using Class C, for example could easily run out of hostid's and need a new Class C address.
- ...need more flexibility within an admin domain

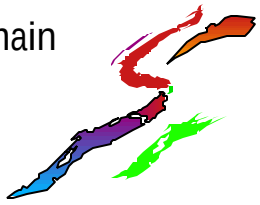

Subnet Addressing

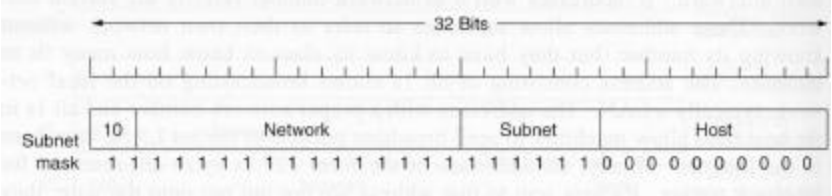

- Instead of considering IP @ as netid+hostid, divide hostid into subnet ID+hostID. (Only known within the netid.)
- NOTE this means that routers must be able to route on subnet ID.
- Above is example BUT subnet/host division could be elsewhere.
  - Often administrators use the natural 8-bit boundary
- Most commonly this is done with Class B addresses.

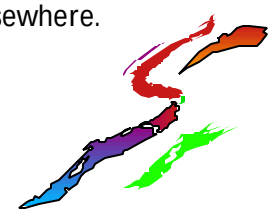

Example Network

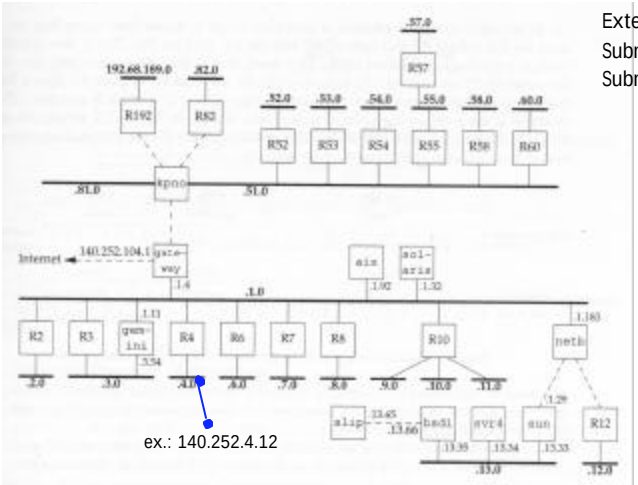

External netid: 140.252 for most.
Subnet mask: 255.255.255.0
Subnets do NOT appear in ext tables.

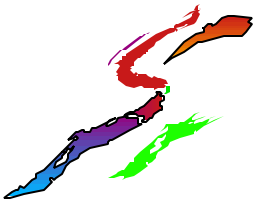

Routing Table Comparisons

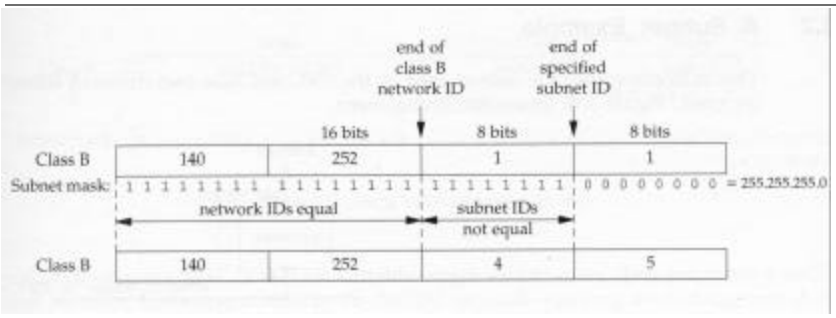

Assume router @ = 140.252.1.1

If dest @ = 140.252.4.5 netid's match and subnets don't.

Use "within netid" table.

If dest @ = 140.252.1.22, subnet's also match, hosts don't.

Route directly to mac @.

If dest @ = 192.43.235.6, netid's don't match.

Route using netid.

Routers need to route only to local hosts, local subnets, and distant network id's.

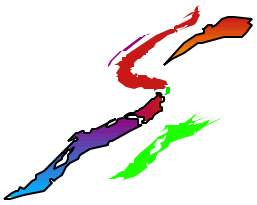

CS & IT Networking Lab (routing config)

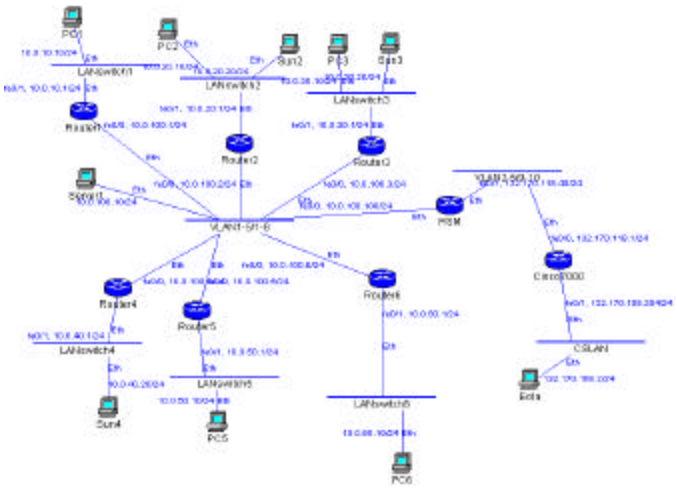

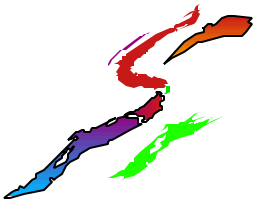

Example Labs

- Lab 1 - Cabling: Types of cabling and how they're wired. Open, crossed, split-pair, NEXT...testing.
- Lab 2 - Introduction to routing (over Ethernet) at workstation. Set IP address and default gateways in Unix and Windows workstations using console. Addess classes and subnet masks. Use of ping, traceroute, telnet, FTP on local segments.
- Lab 3 - CISCO IOS router and 5500 switch configurations. Console port

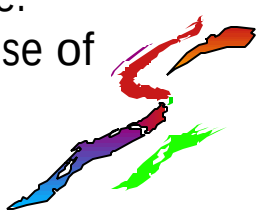

Example Labs *(continued)*

- Lab 4 - Campus connectivity and snoop. Setup VLANS on Route Switch Module. Use of Network Address Translation. Use of snoop on Sun workstations. Use snoop with ping and traceroute. Introduce Netperf software from HP. Use of Snoop and Netperf.
- Lab 5 - Introduction to Net Server. Setting up DNS. Introduction to tftp. Storing and retrieving router and switch configurations. Use of expect and autoexpect.

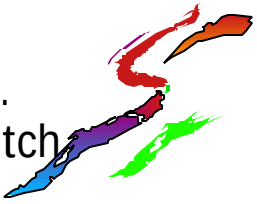

Internet Control Message Protocol

- One of several control protocols: ACMP, ARP, RARP, BOOTP
- Messages:
  - Destination Unreachable
  - Time Exceeded
  - Parameter Problem
  - Source Quench
  - Redirect
  - Echo Request and Echo Reply
  - Timestamp Request and Timestamp Reply

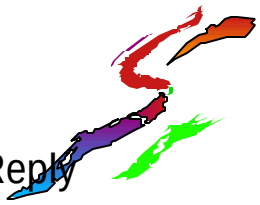

Basic Configuration: Two networks interconnected by a router.

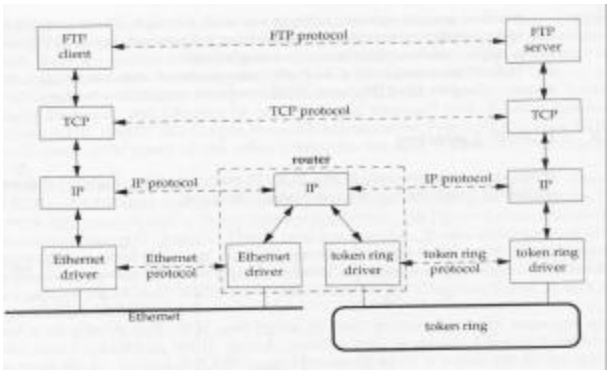

- Ethernet/TR cards send/receive based on 48-bit MAC @.
- NIC's don't look at IP addresses inside IP datagram.
- ARP (Address Resolution Protocol) relates IP/MAC @s.

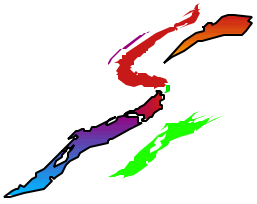

ARP Example (local destination)

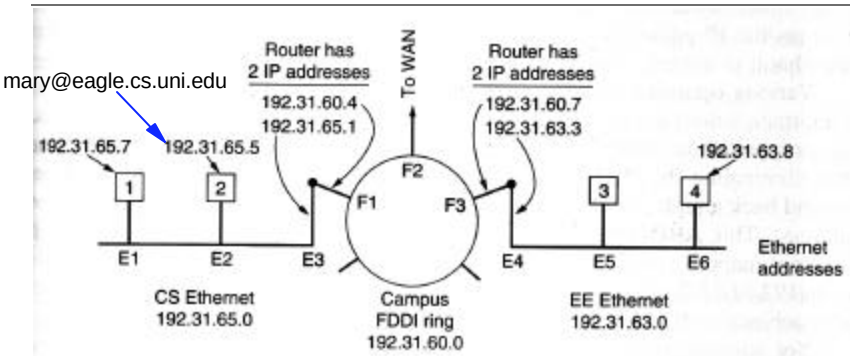

- Host E1 wants to send to mary@eagle.cs.uni.edu
- Lookup done by DNS to get 192.31.65.5.
- E1 broadcasts only on local LAN the ARP request: "Who own's 192.31.65.5?"
- Each NIC will accept the broadcast on MAC and check it's IP address.
- Only E2 responds with its MAC @.
- E1 then puts the original IP datagram inside Ethernet Frame adressed to MAC @ of E2.
- Optimizations: (a) E1 stores info for a time. (b) E1 includes own IP/MAC map in ARP and all stations can copy. (c) Each machine can ARP itself when boots and all stations can copy.

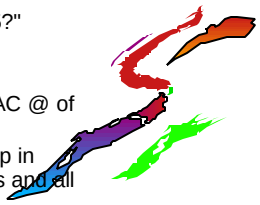

ARP Example (distant destination)

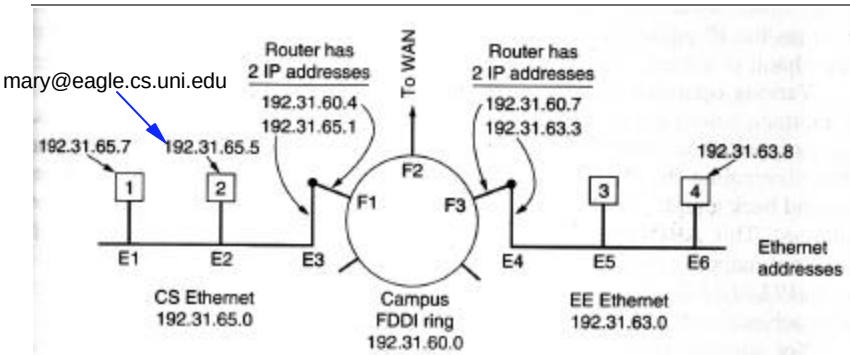

- Host E1 wants to send to E6.
- ARP PROXY:
  - CS Router E3 configured to respond to all ARP requests for 192.31.63.0 with its own IP @ 192.31.65.1.
  - E1 then caches 192.31.65.1/E6.
- Simpler:
  - E1 sees that E6 on distant subnet and it is configured to send all such traffic to E3.
- Either way: IP packet is sent to E3 and E3 looks up the destination IP address.
  - Routing table says send to 192.31.60.7. E3 will ARP on FDDI if necessary to get MAC @.
  - Router E4 receives and looks up destination address in routing table to get hostid. Note that E4 can also send ARP request on the EE Ethernet if necessary to get E6 Mac @.

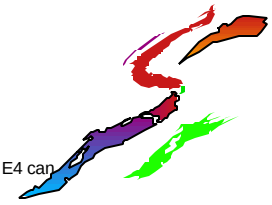

RARP and BOOTP

- ARP gets a Mac @ given an IP @.
- RARP gets IP @ given a Mac @.
- Example: Useful to boot diskless workstation.
  - Identical copies of operating system downloaded to any machine as it boots.
  - Workstation sends broadcast (local LAN) saying "My Ethernet address is e does anyone know my IP address?"
  - RARP server responds (must be on local LAN).
- BOOTP improves on this using UDP msgs forwarded over all routers.
  - Returns IP @ of file server that has memory image, the IP address of default router, and subnet mask.

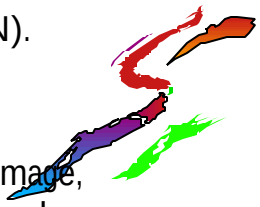

OSPF: "The" Interior Gateway Protocol

- Original Internet IGP was RIP (distance vector) that suffered from count-to-infinity problem.
- Replaced by link-state protocol based on Dijkstra's algorithm in 1979.
- IETF has now "standardized" a new protocol (Open Shortest Path First: RFC 1247).
  - Supports different ToS classes, load balancing, topology changes, hierarchical routing, security.

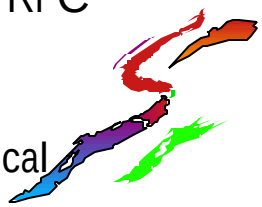

BGP (Border Gateway Protocol)

- BGP is recommended among Autonomous Systems (ASs).
- Referred to as an Exterior Gateway Protocol.
- Support routing based on "policies."
  - No transit traffic through certain ASes.
  - Never put Iraq on route starting from Pentagon.
  - Traffic starting or ending at IBM should not go through Microsoft.
  - ...

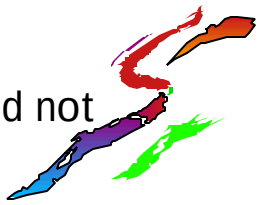

Running Out of Addresses

- Most companies today use Class B addresses: (class A too many hosts, Class C seen as too few, Class B: 65K).
- Class C might have been changed to use 10-bit host numbers instead of 8 (so 1024 hosts per org)...would probably have led to a million or so Class C addresses.
  - a problem for routers because must track all networks
- IETF Invented CIDR (allocate Class C addresses in blocks and by regions) to address this nearterm.

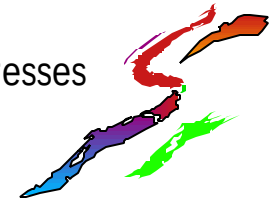

IPV6

- Especially driven by address problems (includes 16 byte addresses).
- Simplifies the header to speed routing
- Includes authentication and privacy features.
- Includes an 8-bit field for ToS (but relatively little progress on how to use it).
- Includes procedures for making V4 to V6 transition.