

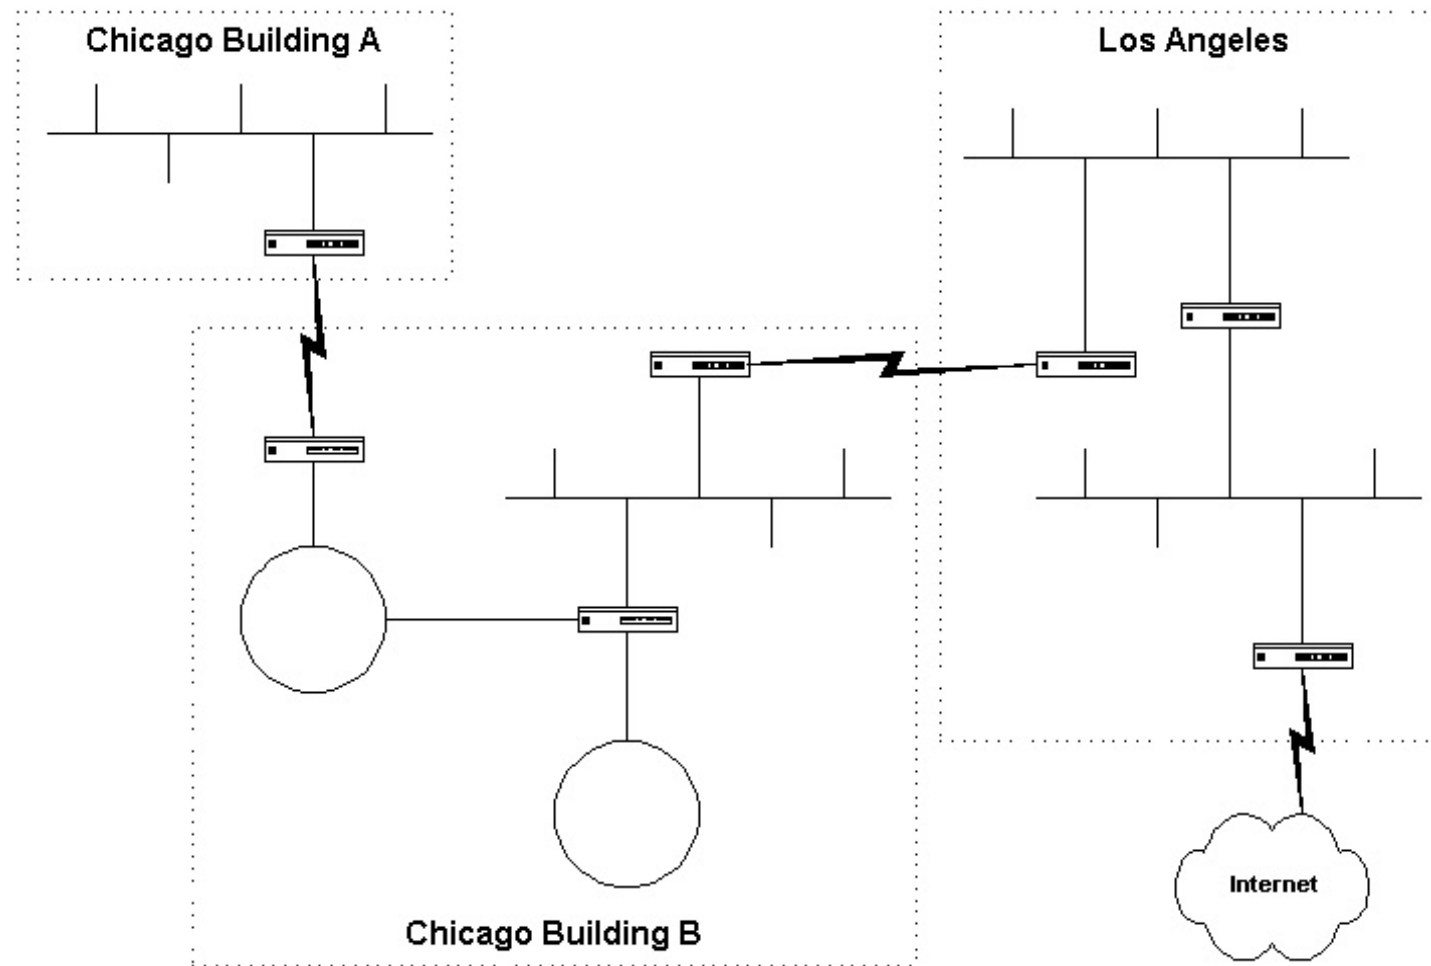
Applied Networks & Security

The Internet Protocol (IP)

<http://condor.depaul.edu/~jkristof/it263/>

John Kristoff
jtk@depaul.edu

Will layer 2 networking suffice?



Layer 2 networking services

- Physical link specific connectivity
- Link unique addressing (identification)
- Limited number of station attachments
- Limitation size, scope and scale

Layer 3 networking services

- Internetworking for data link technologies
- Globally unique addressing
- Scalable (hierarchical) routing
- Common communications format across hosts
- Packet fragmentation capability
- Hardware independent interface
- Packet independence

The Internet Protocol (IP)

- Connectionless
- Unreliable
- Simple (relatively)
- The thin waist in the hourglass model

What can IP do for us?

- Abstracts multiple and various data link networks
- Common communications format
- Hardware independence
- Upper layer independence
- Per-packet independence
- Global, abstracted not data link specific, addressing
- Scalable routing (so far anyway)
- Packet fragmentation capability

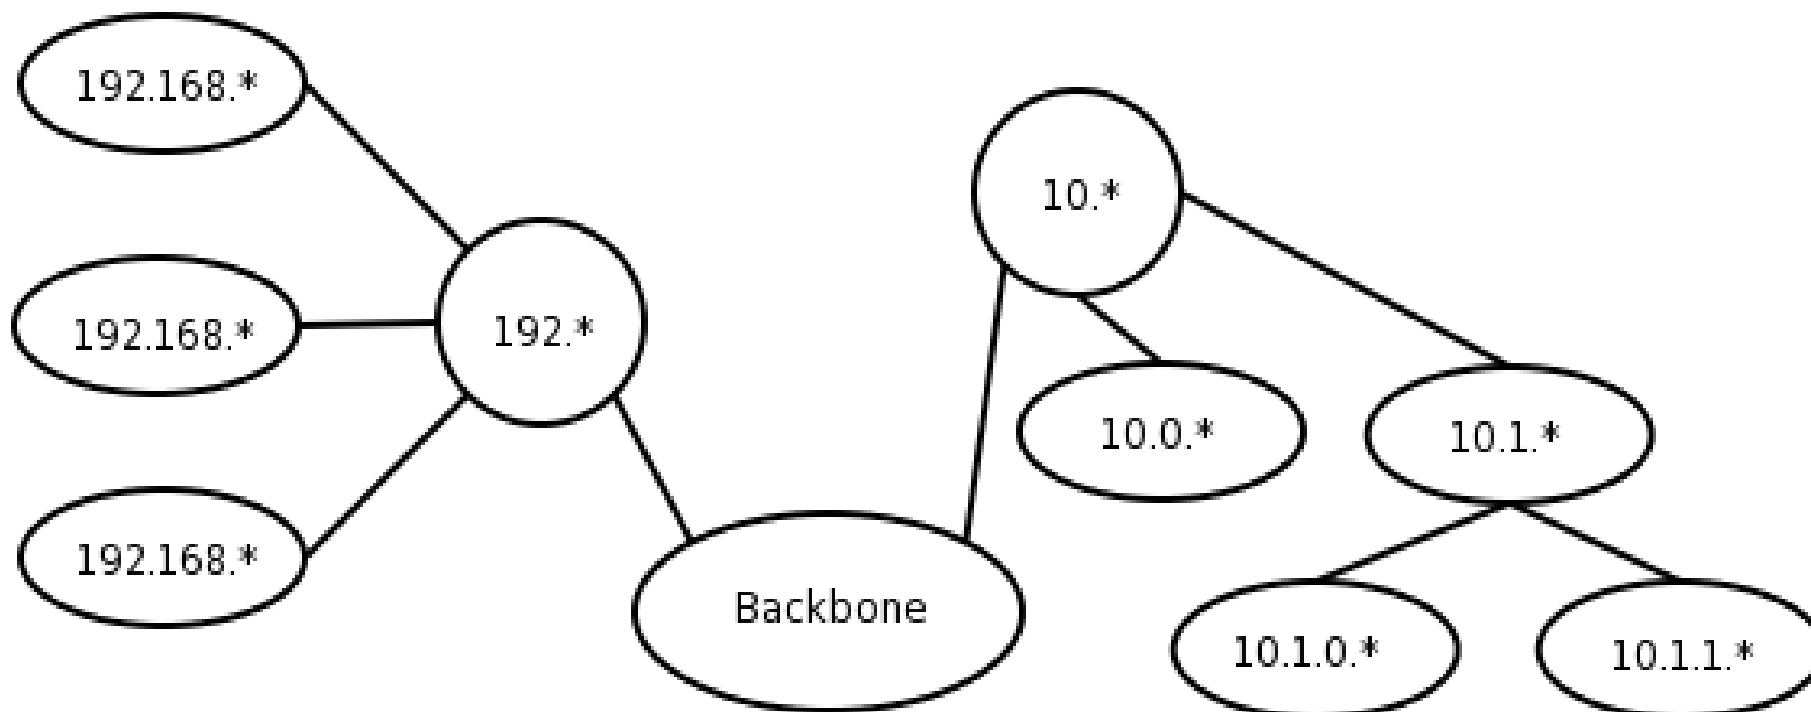
IP layer in perspective

- Layer 1 has repeaters, hubs, modems, etc.
- Layer 2 has bridges/switches
- IP (layer 3) has routers
- Bridges segment layer 2 networks
- Routers segment layer 3 networks
- Hosts need to know nothing about bridges/switches to talk to hosts on the other side of the bridge/switch
- Hosts need to know about routers to talk to hosts on the other side of routers

Bridging versus IP routing

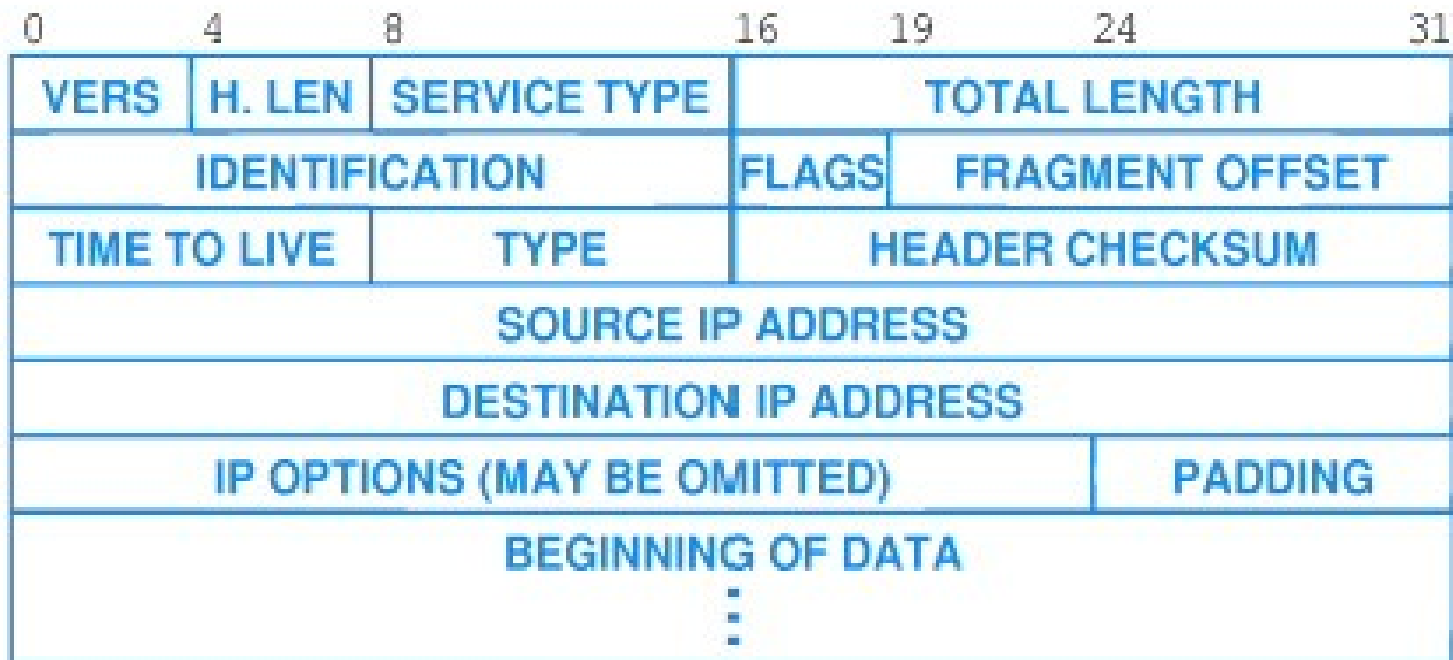
- Bridges learn where all other hosts are by examining source addresses in transmissions
- Bridges learn who is “root bridge” and sets ports to forward or block to avoid loops based on this
- Routers only learn about specific hosts directly attached to local interfaces, using ARP if necessary
- Routers learn (or use static mappings) where other IP networks are and can use one or more equal paths to “forward” packets to them

Hierarchical Routing



IP datagram

*diagram courtesy of <http://www.netbook.cs.purdue.edu>



Inside an IP datagram

- Version field
 - Usually set to binary 0100 (is what in decimal?)
- Header length
 - Length of IP header in 32-bit words (4 octets)
 - Typically set to 5 (as in $5 * 4$ octets = 20 bytes)
- Type of Service (Tos) – redefined in newer RFCs
 - An indication of quality/class of service
 - Rarely used with success outside a single AS

Inside an IP datagram [cont.]

- Total length
 - total IP datagram length in octets
 - maximum value is 65535, but rarely > 1500
- Identification
 - to identify fragments of a single IP datagram
 - experimentally used in tracing DDoS sources
- Flags
 - bit 0 reserved
 - others for fragmentation handling (DF/MF)

Inside an IP datagram [cont.]

- Fragment offset
 - helps piece together fragments
- Time to live (TTL)
 - limits the number of router hops datagram incurs
 - counts down to zero, at zero it is discarded
- Protocol type
 - indicates next (upper?) layer protocol in payload
 - Does it have to be an “upper” layer?

Inside an IP datagram [cont.]

- Header checksum
 - used to verify header validity at each hop
- Source/Destination address
 - 32-bit addresses
- Options (optional, duh)
 - rarely used, padded to 32-bit boundary if needed
- Payload (next protocol plus it's data)
 - variable length

IP address

- Virtual, not specific to a hardware device
- 32-bit fixed address length (IPv4)
- Unique address for each interface (typically)
- Global registrar or upstream provider assigns network bits (prefix)
- Local network admin assigns subnet and host bits (suffix)
- Usually written in dotted decimal (dotted quad)
 - e.g. 140.192.5.1

IP address types

- Unicast (one-to-one)
 - source addresses should always be unicast
- Multicast (one-to-many)
 - receivers join/listen to group destination address
- Broadcast (one-to-all)
 - special case of multicast, usually unnecessary
- Anycast (one-to-one-of-many)
 - usually one-to-nearest, often used for reliability

IP address notation

*diagram courtesy of <http://www.netbook.cs.purdue.edu>

<u>32-bit Binary Number</u>	<u>Equivalent Dotted Decimal</u>
10000001 00110100 00000110 00000000	129 . 52 . 6 . 0
11000000 00000101 00110000 00000011	192 . 5 . 48 . 3
00001010 00000010 00000000 00100101	10 . 2 . 0 . 37
10000000 00001010 00000010 00000011	128 . 10 . 2 . 3
10000000 10000000 11111111 00000000	128 . 128 . 255 . 0

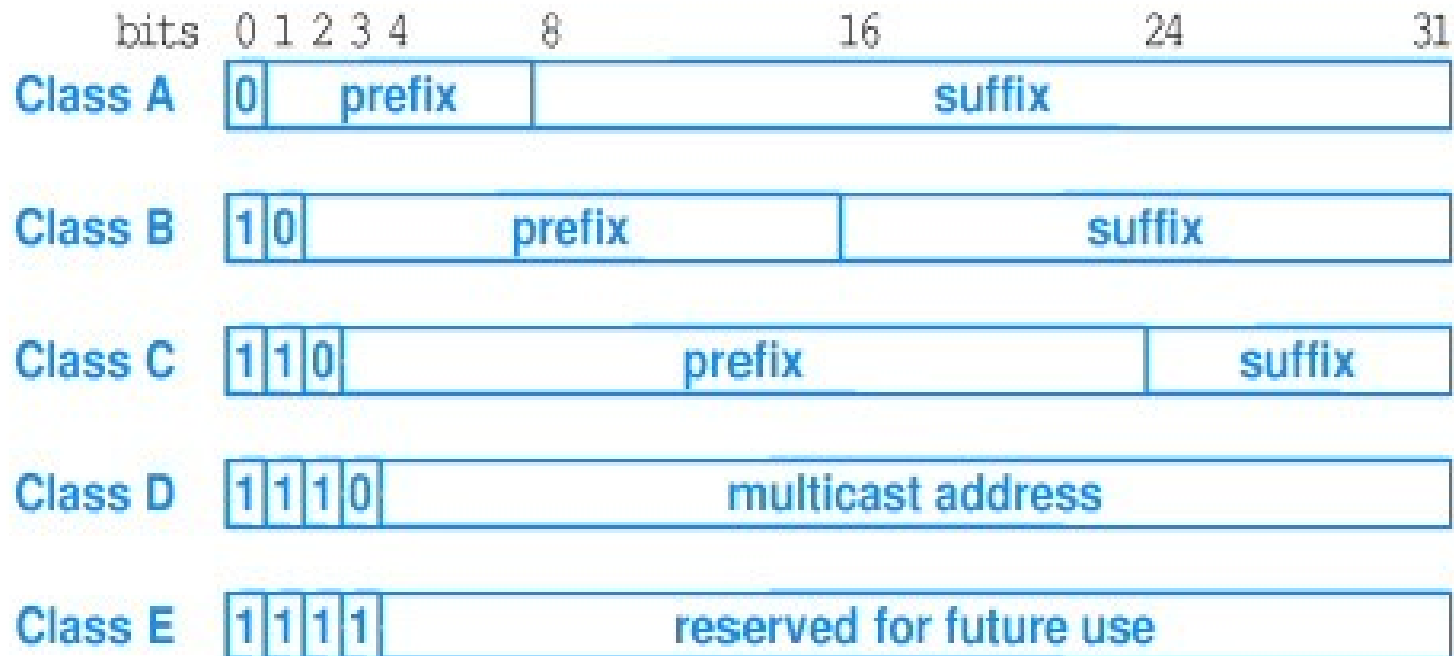
Special IP addresses

*diagram courtesy of <http://www.netbook.cs.purdue.edu>

Prefix	Suffix	Type Of Address	Purpose
all-0s	all-0s	this computer	used during bootstrap
network	all-0s	network	identifies a network
network	all-1s	directed broadcast	broadcast on specified net
all-1s	all-1s	limited broadcast	broadcast on local net
127	any	loopback	testing

Classful IP addressing

*diagram courtesy of <http://www.netbook.cs.purdue.edu>




Classful address sizes

*diagram courtesy of <http://www.netbook.cs.purdue.edu>

Address Class	Bits In Prefix	Maximum Number of Networks	Bits In Suffix	Maximum Number Of Hosts Per Network
A	7	128	24	16777216
B	14	16384	16	65536
C	21	2097152	8	256


Example IP network

*diagram courtesy of <http://www.netbook.cs.purdue.edu>



Example IP router addressing

diagram courtesy of <http://www.netbook.cs.purdue.edu>



Let's look at a “route server”

- <http://www.routeviews.org>
- We'll telnet into a router and look around, particularly at the routing table


Classful addressing limitations

- Internet growth and address depletion
- Route table size (potentially lots of class C nets)
- Misappropriation of addresses
- Lack of support for varying sized networks
 - Class B is often too big, Class C often too small

IP addressing solutions

- Subnetting
- Supernetting
- Classless interdomain routing (CIDR)
- Variable length subnet masks (VLSM)
- BOOTP and DHCP (temporary addresses)
- NATs with port address translation (yucky!)

Subnetting



Subnet masks

- The bit length of the prefix (network bits)
- Prefix (network bits) no longer classful (fixed size)
- Use of the slash '/' notation to represent addresses
 - 140.192.5.1 with mask of 255.255.255.128 is:
 - 140.192.5.1/25
- As viewed in binary for clarity, a /25 mask is:
 - 11111111.11111111.11111111.10000000

Subnet masks example

- Given 140.192.50.8/20 what is the...
 - subnet mask in dotted decimal notation?
 - directed broadcast address in dotted quad?
 - total number of hosts that can be addressed?

Supernetting

- Combine smaller address blocks into an aggregate
- If class B is too big and class C is too small...
 - Combine 199.63.0.0/24 to 199.63.15.0/24
 - To form 199.63.0.0/20

Supernetting example

- Given an ISP that has 128.15.0.0/16:
 - what block might be assigned to a customer needing to address 300 hosts?
 - how does the ISP manage their IP address allocation if there are many customers with varying address requirements?

CIDR

- Routers using aggregated prefixes (CIDR blocks)
 - primarily through the use of supernetting
- So instead of routing multiple class C blocks...
 - ...advertise some larger aggregate, e.g. /20
- The Internet CIDR report:
 - <http://www.cidr-report.org>

CIDR example

- Given an ISP that announces:
 - 64.5.0.0/20
 - 64.5.16.0/20
 - 192.0.2.0/25
 - 192.0.2.192/26
 - 192.0.2.128/26
- What is the least number of CIDR announcements that can be made for this ISP?
- Why might address blocks not be aggregated?

VLSM

- Many subnet sizes in an autonomous system (AS)
- Allows for efficient use of address space
- Can be used to build an internal hierarchy
- External view of the AS does not change
- An AS may be allocated 140.192.0.0/16, but...
- internally may use:
 - 140.192.0.0/17
 - 140.192.128.0/24
 - 140.192.129.0/25 and so on...

VLSM example

- Given an assignment of 140.192.0.0/16, create an addressing strategy to support:
 - 6 satellite offices and 1 large headquarter site
 - 6000 total hosts on all combined networks
 - headquarters needs about 50% of all addresses
 - satellite offices need 200 to 700 addresses
 - overall growth per year is 500 hosts

Obtaining IP addresses

- IANA has global authority for assignment
- Regional Internet Registries (RIRs) delegate directly to ISPs and sufficiently large nets
- ISPs assign addresses to end users and smaller networks
- RFC 1918 defines private address blocks
 - NOT globally unique
 - NOT for hosts attached directly to public Internet
 - 10.0.0.0/8, 172.16.0.0/12 and 192.168.0.0/16

Steve Deering Hourglass Video

- IETF 51
- <http://www.iab.org/documents/docs/hourglass-london-ietf.pdf>
- <ftp://videolab.uoregon.edu/pub/videolab/video/ietf51/ch2/>