SC/CSE 3213 Winter 2014

L3: Internet Structure and Addressing

Sebastian Magierowski
York University

Outline

- Basic structure
- Addressing
- Subnetworks
- Supernetworks
- DNS

Internet: An Internetwork

- A collection of many different networks
  - millions?
- The big ones are called Autonomous Systems (AS)
  - ISPs, campus networks, government, etc.
    - York
    - Bell
    - AT&T
    - Global Crossing
  - From a few to 1000s of routers
  - Typical path goes through < 6 ASes

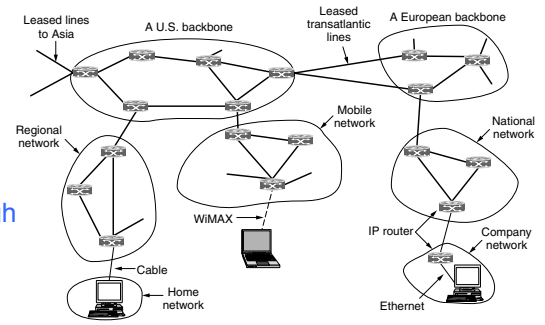

[©Pearson, Tanenbaum]

Autonomous Systems (Domains)

- Almost 50,000 and growing [bgp.potaroo.net]
  - further subdivided into…
    - "areas"
    - improves routing efficiency

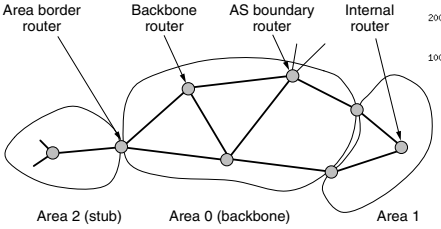

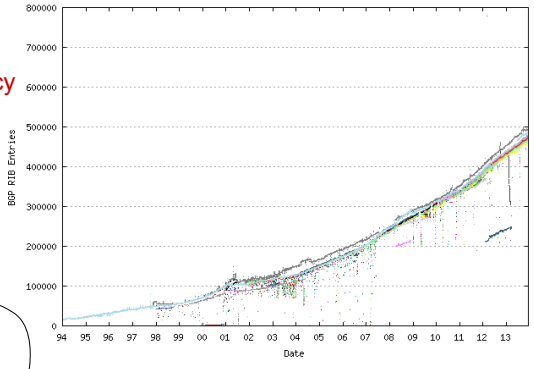

[©Pearson, Tanenbaum]

Peering

- ASes connect at IXPs (Internet eXchange Points)
  - room full of routers at least one per ISP
  - routers connected by LAN
  - e.g. AMS-IX (in Amsterdam) connects hundreds of ISPs
    - Transit AS
    - Stub AS
    - Multihomed AS

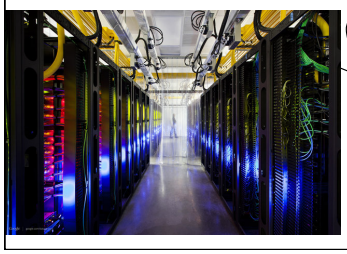

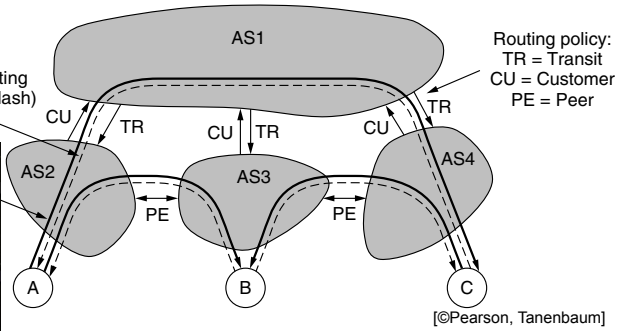

[©Pearson, Tanenbaum]

TorIX

- 151 Front Street
  - 21st or so
    - 83 Gbps on avg.
    - 175 members
  - AMS-IX
    - 1571 Gbps
    - 613 members

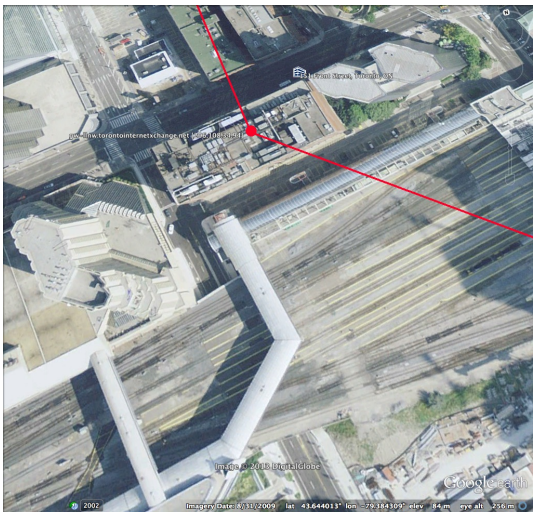

Message Switching

- Your signals route through these networks [www.ixmaps.ca]
  - T.O. to hawaii.edu
    - 14 hops
    - to Australia before Hawaii?
    - through a suspected NSA listening facility
- Messages are are broken into small parts called packets (datagrams)
- Any of a number of paths may be routed as part of packet switching

Packets

- A piece of your digital message
- Why would you want to break your message up?
- Split into a…
  - Payload
    - The data you want to send
    - Typically less than 64 KB
    - And often < 1.5 KB
  - Header
    - Information about what you want the network to do with the packet
    - Size variable, roughly 20-40 bytes
- Most important network job?
  - Deliver the packet!

01HEADER11 1011101PAYLOAD100010

Addressing

- How do packets get where they are supposed to go?
- Header contains source and destination address
  - Each connection (link interface to host or router) is given a unique address
  - Each router contains a file (often stored in specialized hardware) called a forwarding table telling the packet where to hop next

Destination	Interface

Addressing Hierarchy

- *M* networks with *N* hosts each require *M•N*
- To simplify the table, a hierarchical address structure is used
  - net ID + host ID
  - net ID also called the prefix
- Only M entries needed
  - Instead of needing billions of entries Internet routers can work with "only" millions ☹
  - More on this in a moment (CIDR & BGP)

Destination	Next Hop

IP Addresses

- 32-bit string
  - 10000000.11010000.00000011.00010011
- Written one byte at a time in dotted decimal notation
  - 128.208.3.19
- IPv4
- IPv6: 128-bit string

Where do IP Addresses Come From?

- ICANN
  - Internet Corporation for Assigned Names and Numbers
  - non-profit in Los Angeles, USA
  - incorporated in 1998 (U.S. Govt. was previous manager)
  - IANA (Internet Assigned Numbers Authority)
- Delegates to Regional Internet Registries (through IANA)
  - Organized into NRO (Number Resource Organization)
  - ARIN for North Am. & Carribean
    See: https://www.arin.net/knowledge/ip_blocks.html for list of IP address ranges

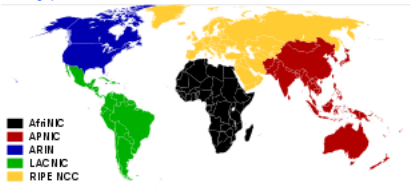

IP Addresses with Class

- Address used to come in well defined classes
  - Class A (8-bit prefix)
    - Net ID starts with 0_2
    - 2^7-1 = 127 networks of 16M hosts
  - Class B (16-bit prefix)
    - Net ID starts with 10_2
    - 16k nets of 64k hosts
  - Class C (24-bit prefix)
    - Net ID starts with 1102
    - 4M nets of 255 hosts

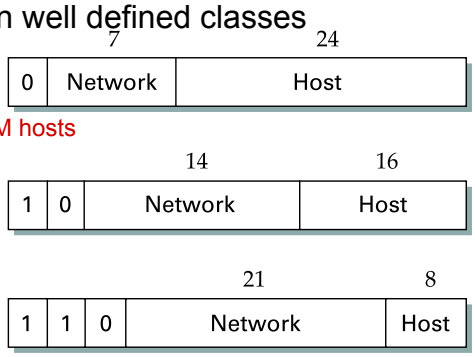

- Net ID of an address relatively easy to identify
  - just look at first few bits to determine format of Net ID

[©Elsevier, Peterson/Davie]

Classfull Address Problems

- Three basic classes waste too many addresses

Subnets

- Use a finer partitioning scheme in your network
  - Give each of your subnetworks a Subnet ID (take some digits from Host ID)
- For example:
  - 10000000.11010000.XXXXXXXX.XXXXXXXX
- Can become…
  - 10000000.11010000.1XXXXXXX.XXXXXXXX
    - 17-bit prefix size
  - 10000000.11010000.00XXXXXX.XXXXXXXX
    - 18-bit prefix size
  - 10000000.11010000.011XXXXX.XXXXXXXX
    - 19-bit prefix size

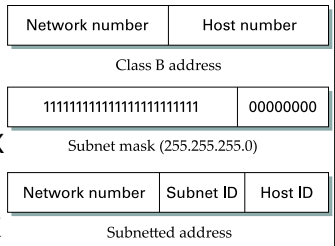

Subnet Masks

- Each subnetwork is *identified* with a subnet mask
  - 32-bit number with 1's in the prefix positions and 0's elsewhere
- For example…
  - 10000000.11010000.1XXXXXXX.XXXXXXXX
    - 11111111.11111111.10000000.00000000
    - Subnet mask: 255.255.128.0
  - 10000000.11010000.00XXXXXX.XXXXXXXX
    - 11111111.11111111.11000000.00000000
    - Subnet mask: 255.255.192.0
  - 10000000.11010000.011XXXXX.XXXXXXXX
    - 11111111.11111111.11100000.00000000
    - Subnet mask: 255.255.224.0

Subnet Addresses

- Subnet address is…
  - AND of any address in subnet with its mask
  - Use /# to denote number of network (+ subnet) bits (CIDR notation)
- For example…
  - 10000000.11010000.1XXXXXXX.XXXXXXXX
    - Subnet mask: 255.255.128.0
    - Subnet address: 128.208.128.0/17
  - 10000000.11010000.00XXXXXX.XXXXXXXX
    - Subnet mask: 255.255.192.0
    - Subnet address: 128.208.0.0/18
  - 10000000.11010000.011XXXXX.XXXXXXXX
    - Subnet mask: 255.255.224.0
    - Subnet address: 128.208.96.0/19

Subnet Forwarding

- When router gets any address it ANDs it with possible subnet masks
  - 255.255.128.0 (CS)
  - 255.255.192.0 (EE)
  - 255.255.224.0 (Art)

Destination	Interface
128.208.128.0/17	CS
128.208.0.0/18	EE
128.208.96.0/19	Art

- And compares result to prefix bits in forwarding table
- E.g.: 128.208.2.15

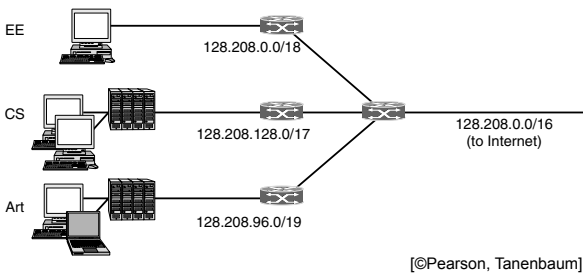

[©Pearson, Tanenbaum]

More Problems with Address Classes

- Subnetting helps optimal address utilization
- But…
  - It still leaves us with a lot of networks to find
  - Huge forwarding and routing tables are required to deal with all possibilities
- Technically it still works within the old class-based system (i.e. just splits up coarse classes)
- If we do away with the classfull system altogether we can conceptually aggregate subnetworks into larger units
- This slows table growth

Supernetting/Aggregation/CIDR

- Classless Interdomain Routing (CIDR)
  - A set of rules allowing routers to arrange network addresses into single blocks

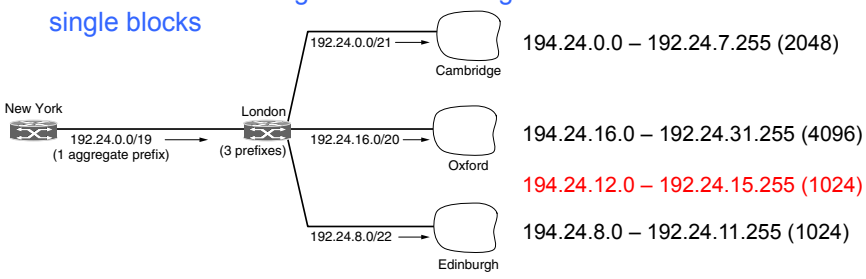

  - After observing traffic patterns a router automatically aggregates the 3 nets into a single address
    - 192.24.0.0/19

[©Pearson, Tanenbaum]

Supernets: Longest Matching Prefix

- Prefixes can overlap
  - Routers send information along the longest matching prefix

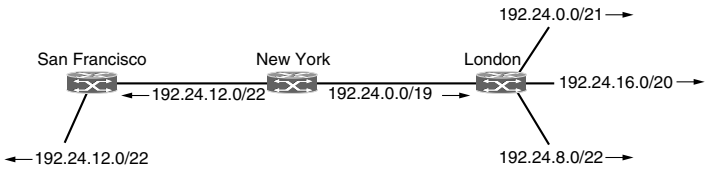

- Classless addressing has replaced classfull addressing
  - more efficient use of addresses
  - smaller forwarding/routing tables (~200,000 entries in core routers rather than millions)

[©Pearson, Tanenbaum]

DNS

- We (people) don't use IP addresses
  - www.yorku.ca
- How do host names get converted to numerical addresses?
  - Too many addresses for small host computers to map
  - Conflicts would occur unless map was centrally managed
- Domain Name System (DNS)
  - A hierarchical naming scheme and distributed database system
  - library procedure called a resolver

Hierarchical Namespace

- Like a file tree with names separated by dots rather than /: eecs.yorku.ca
- Over 250 top-level domains
  - Managed by ICANN
- Second-level domains available from ICANN registrars

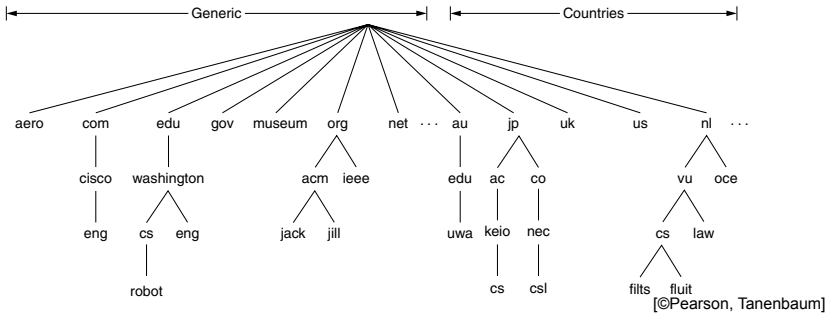

[©Pearson, Tanenbaum]

Zones

- Hierarchy partitioned into subtrees called zones
- Each zone is under the control of some administrative authority
  - Implemented with two or more name servers

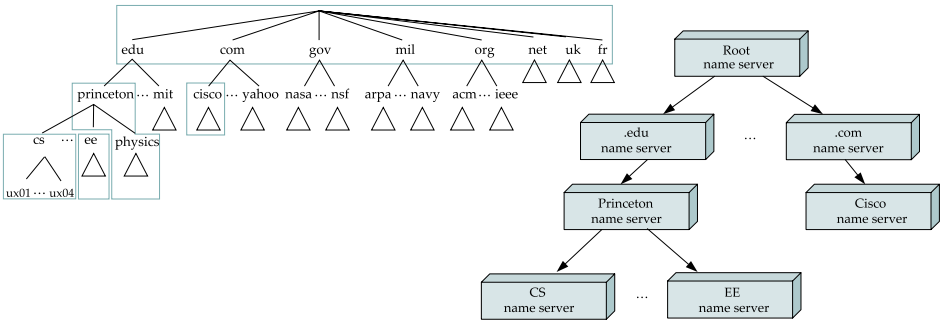

Name Resolution

- A combination or recursive and iterative queries

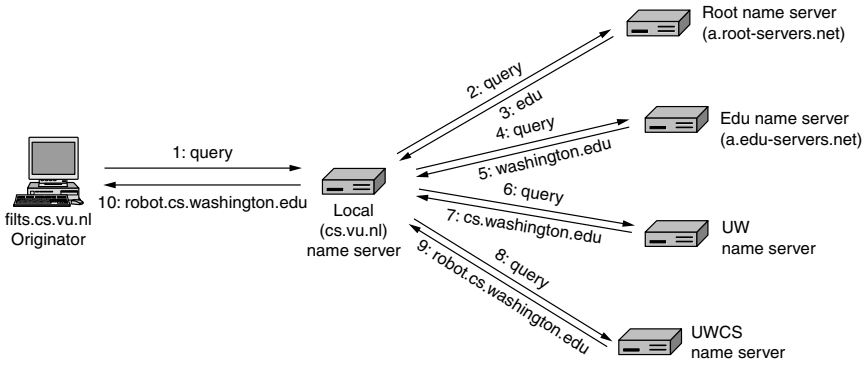

[©Pearson, Tanenbaum]

Name Server Trivia

- 13 root DNS servers
  - a-root-servers.net…m-root-servers.net
  - actually a number of geographically distributed machines (with same IP address)
  - reached by anycast routing
    - a routing method delivering to nearest member of a group
    - group looks like one logical node