CS 470
Spring 2021

Mike Lam, Professor

Naming

Content taken from the following:

"Distributed Systems: Principles and Paradigms" by Andrew S. Tanenbaum and Maarten Van Steen (Chapter 4)
Various online sources (including openclipart.org)

Naming

- "What's in a name?"
  - *"That which we call a .com by any other TLD would load just as quickly."*

"There are only two hard things in Computer Science: cache invalidation and naming things."

- Phil Karlton (Netscape)

Trivia

- What is *Netscape*?
  - A. A web browser
  - B. A web directory
  - C. An internet service provider
  - D. A brand name
  - E. All of the above
  - F. None of the above

Addressing

- Concept of an entity and its name vs. its address
- Some names are true identifiers
  - Each identifier refers to at most one entity
  - Each entity is referred to by at most one identifier
  - Identifiers are never re-used at another time
- Name-to-address binding
  - Name space: domain of all possible names
  - Static vs. dynamic
  - Central vs. decentralized
    - Name server: central host responsible for maintaining bindings

Addressing

- Which of the following is the most **decentralized** name binding?
  - A. Mailing addresses
  - B. Socrative participant names
  - C. Subreddits
  - D. Human nicknames
  - E. Xbox gamertags

Naming schemes

```
eb40af8e
c6c1904c
0eceda3e
28dec8ba
4b6683e7
88c9618b
3566223f
38b22b10
```

Flat

```
/
  bin/
    bash
  etc/
    passwd
  usr/
    bin/
      nano
      vim
    lib/
```

Structured

```
(444, Molloy)
(445, Sprague)
(455, Aboutabl)
(458, Heydari)
(470, Lam)
(482, Wang)
(482, Tjaden)
(488, Johnson)
```

Attribute-based

Flat naming

- Identifiers contain no location information
- Various lookup approaches
  - Broadcast / multicast
  - Forwarding pointers
  - Proximity routing
- Examples: ARP, Chord

Distributed hash tables

- Chord uses an m-bit identifier space and modulo arithmetic
  - Key k is stored at the node with the smallest id ≥ k
- Each node maintains a finger table of forward shortcuts
- To look up k, repeatedly follow lookups in finger table
  - Goal: halve distance to destination every hop

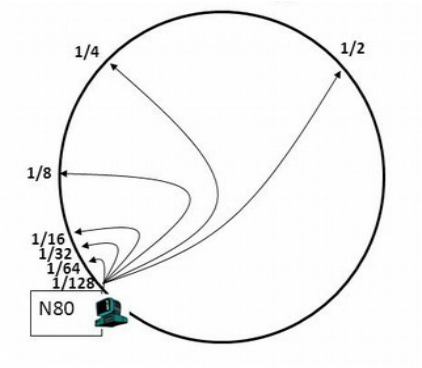

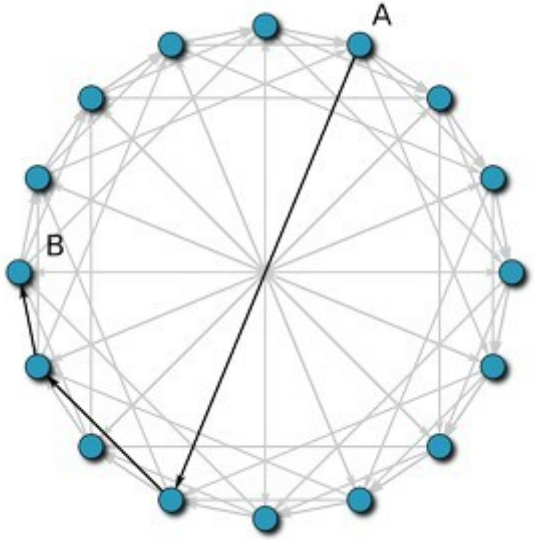

Addressing

- Which of the following is the maximum size of the finger table for a 256-node Chord network?
  - A. 0
  - B. 1
  - C. 8
  - D. 32
  - E. 128

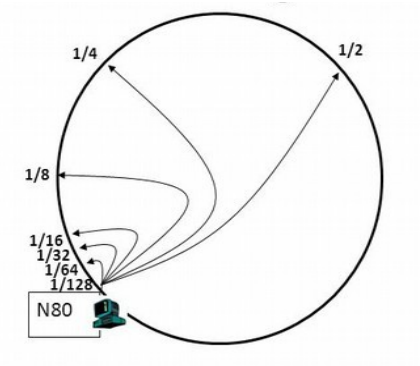

Structured naming

- Root vs. interior vs. leaf nodes
- Absolute vs. relative names
  - Global vs. local names
- Iterative vs. recursive resolution
- Linking and aliasing
  - Hard vs. soft (symbolic) links
- Mounting and mount points
- Examples: file systems, DNS, NFS

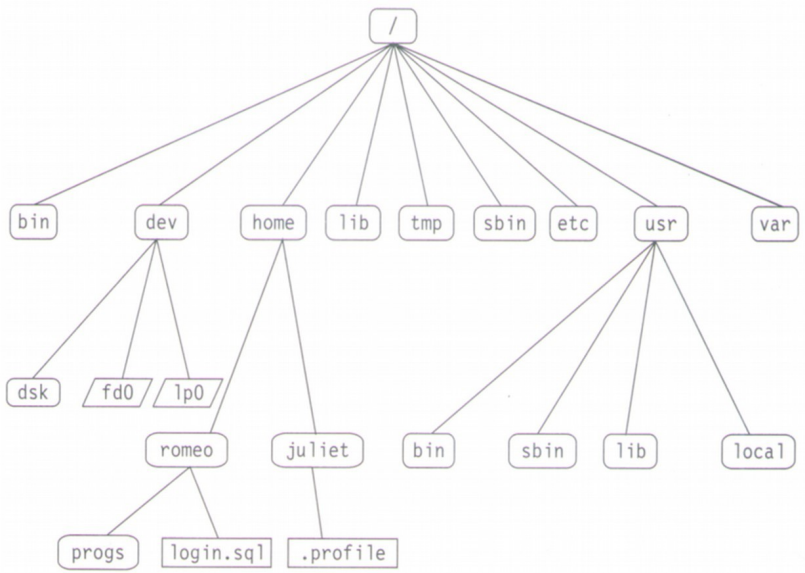

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/mapper/rhel_login01-root	50G	23G	28G	46%	/
/dev/sda6	497M	206M	292M	42%	/boot
nfs.cluster.cs.jmu.edu:/nfs/home	100G	4.6G	96G	5%	/nfs/home
nfs.cluster.cs.jmu.edu:/nfs/scratch	2.0T	862G	1.2T	43%	/scratch

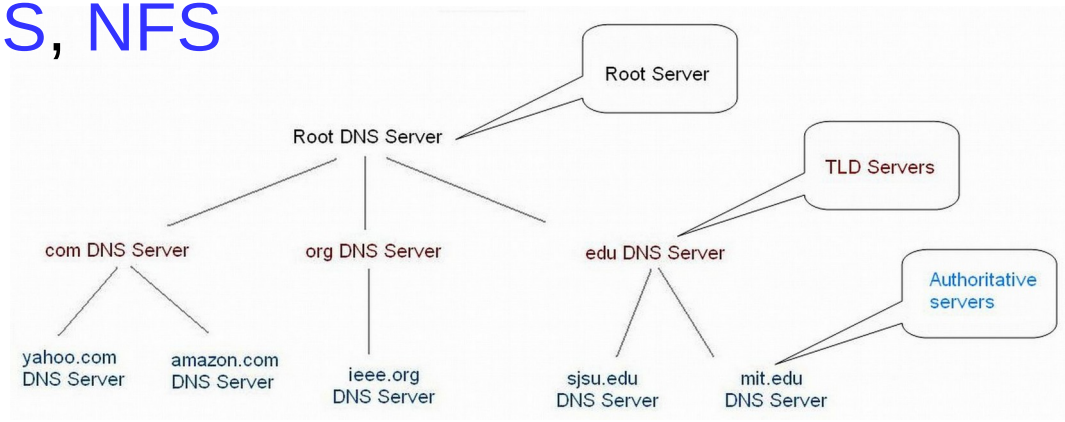

Naming

- Which of the following is an example of a **structured** (as opposed to **flat**) name binding?
  - A. Mailing addresses
  - B. Socrative participant names
  - C. Subreddits
  - D. Human nicknames
  - E. Xbox gamertags

IPv4

- IPv4: 32 bits - four octets w/ CIDR notation (/8, /16, etc.)
  - Classful addressing: Class A, Class B, Class C
  - IETF and IANA allocate addresses (32 bits - 4 billion total addresses)
  - Published in 1981; now nearly exhausted
- Notable networks
  - Private (10.0.0.0/8)
  - Loopback (127.0.0.0/8)
  - JMU (134.126.0.0/16)
  - Private (192.168.0.0/16)

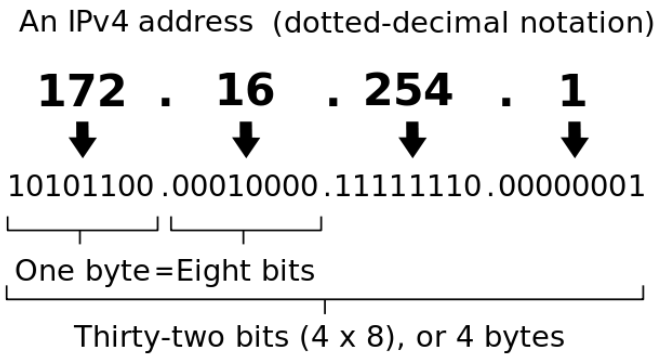

from https://en.wikipedia.org/wiki/IPv4

https://xkcd.com/195/

IPv4 map

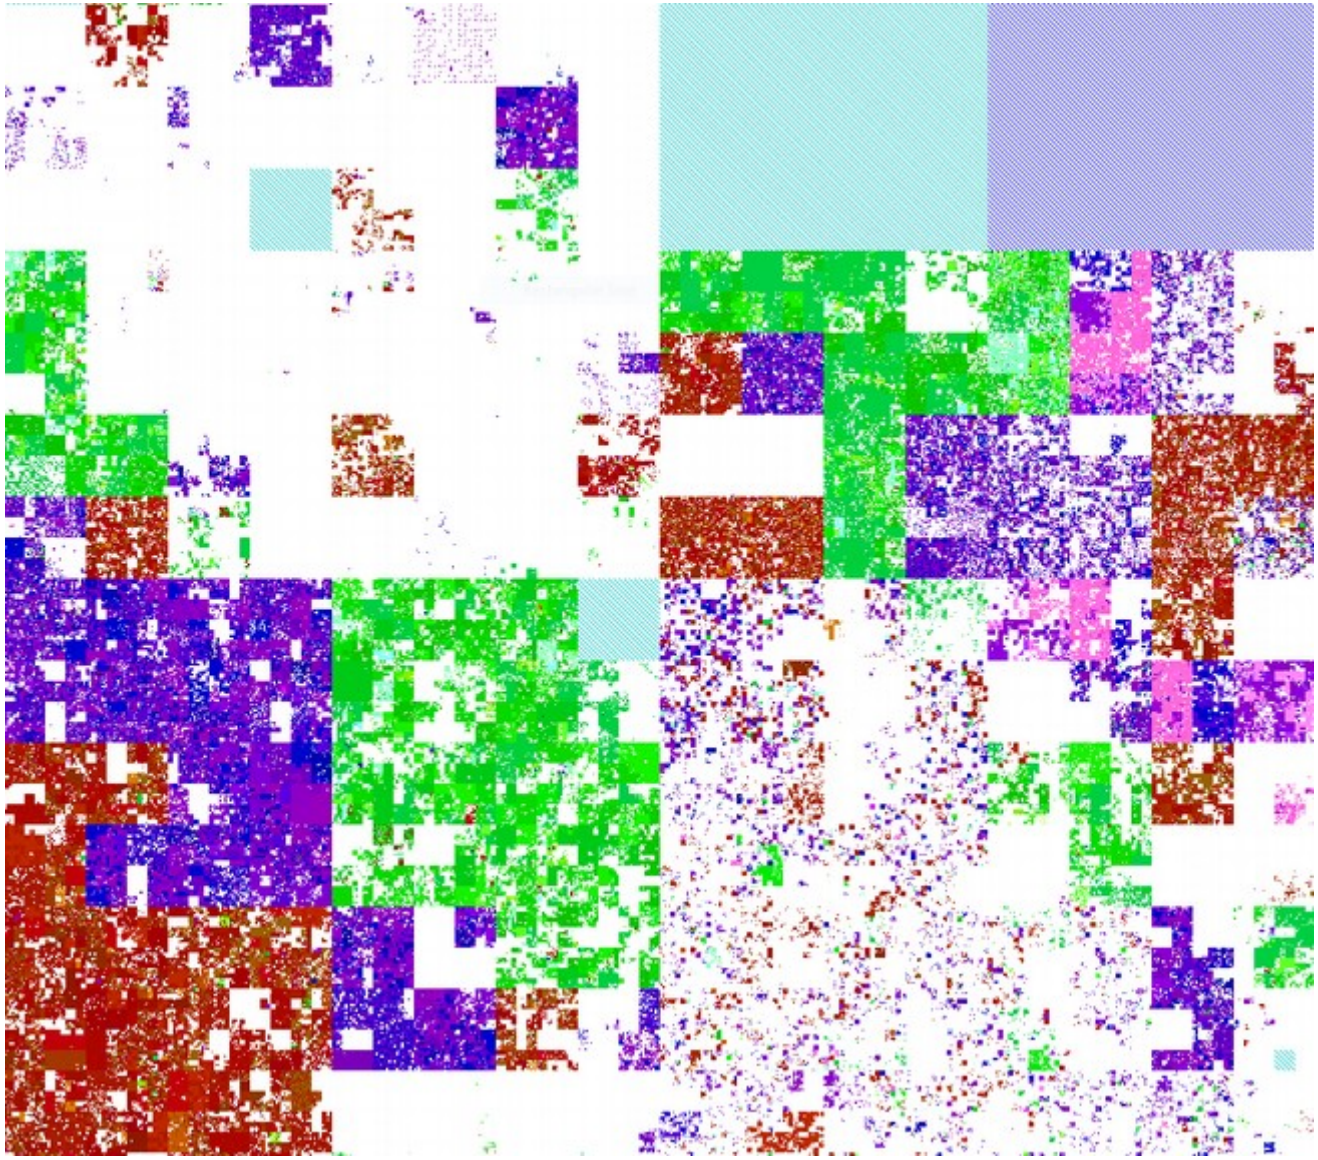

from `https://ant.isi.edu/address/browse/index.html`

IPv4

- What is the total number of addresses in IPv4?
  - A. 2^8
  - B. 2^{16}
  - C. 2^{32}
  - D. 2^{64}
  - E. 2^{128}

IPv6

- IPv6 published in 1998
  - 128 bits - 3.4×10^{38} total addresses
  - Eight groups of 16 bits (4 hex chars)
  - 64-bit routing prefix, 64-bit host/interface identifier

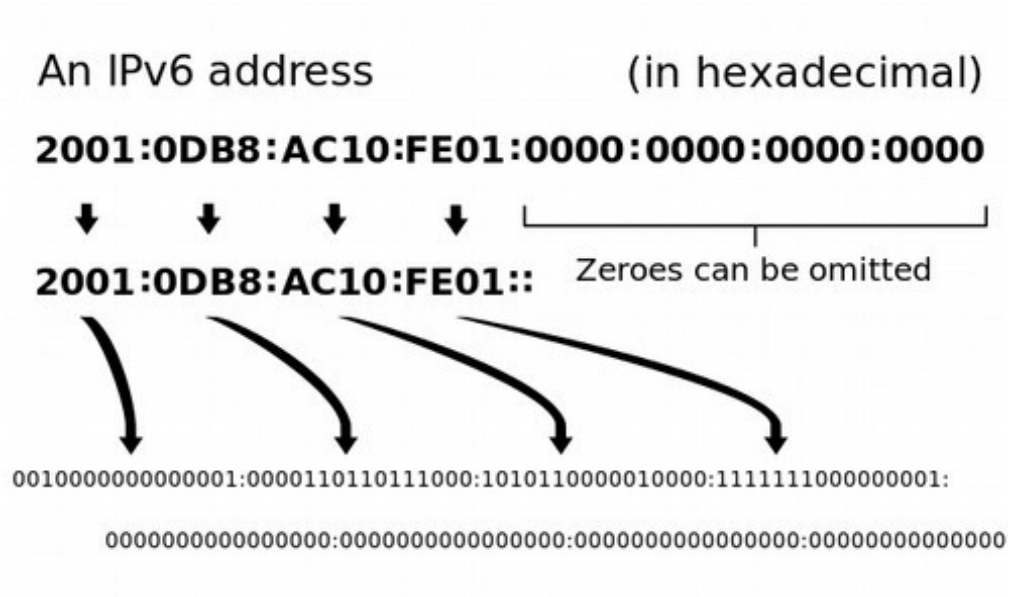

from `https://en.wikipedia.org/wiki/IPv6`

IPv6

- What is the total number of addresses in IPv6?
  - A. 2^8
  - B. 2^{16}
  - C. 2^{32}
  - D. 2^{64}
  - E. 2^{128}

IPv4 vs. IPv6

- The IPv6 name space is far larger than you think!
  - In fact, there is NO WAY to draw the two address spaces to scale. If IPv4 were a 1.6-inch square, IPv6 would be a square the size of the solar system!
  - $2^{128} \approx 10^{38} \gg$ the number of drops of water in all the world's oceans (10^{25}) or the number of stars in the observable universe (10^{23})
  - "If we had been assigning IPv6 addresses at a rate of 1 billion per second since the earth was formed, we would have by now used up less than one trillionth of the address space."
  - "We could assign an IPv6 address to every atom on the surface of the earth – and have enough addresses left over for another hundred earths."

Sources:

- http://waitbutwhy.com/2014/11/1000000-grahams-number.html
- http://www.tcpipguide.com/free/t_IPv6AddressSizeandAddressSpace-2.htm
- http://www.brucebnews.com/2010/10/ipv6-and-really-large-numbers/

Why haven’t we transitioned?

- Advantages
  - Solves IP naming problem pseudo-permanently
  - Internet of Things (IoT) threatens to explode the number of devices requiring an address
  - Increasing cost to acquire IPv4 addresses
- Obstacles
  - Network Address Translation (NAT) allows multiple hosts to use a single public IP address
  - IPv4 blocks have become more “fluid”
  - Lack of expertise managing multi-protocol networks

Informed by discussion with Mike Ripley, JMU Information Technology

Attribute-based naming

- Human-friendly resource identifiers
- Storage of (key, value) pairs
- Often implemented with distributed hash tables
  - Centralized vs. decentralized lookups
  - You will implement this in P4!
- Semantic overlay networks
  - Nodes maintain explicit links to "semantically proximate" nodes
  - Most useful in distributed peer-to-peer networks
  - Exploit small-world effect

Attribute-based naming

- Which of the following is the best example of a **semantic overlay** network?
  - A. Mailing addresses
  - B. Socrative participant names
  - C. Subreddits
  - D. Human nicknames
  - E. Xbox gamertags